UNIVERSITY of CAMBRIDGE

PROCEEDINGS OF NETWORKSHOP 3

27 - 28 SEPTEMBER 1978

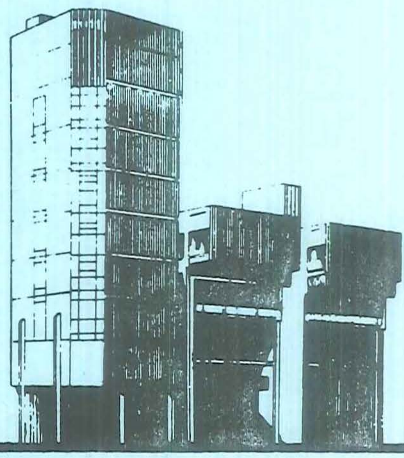

University of Cambridge
Computer Laboratory
Corn Exchange Street
Cambridge CB2 3QG
England

CONTENTS

INTRODUCTION

The papers in this document record the proceedings of the Networkshop held at the University of Cambridge on 27 - 28 September 1978. This was the third in a series of inter-university workshops organised by the Network Unit of the Computer Board and Research Councils, and was attended by representatives from almost every British university and from certain other institutions.

The main themes of this Networkshop were:

(a) Post Office Packet Switched Service (PSS)
- latest details supplied by the Post Office
- University and SRC requirements

(b) Network services
- to support different types of system
- for campus transmission and switching
- implementation on mainframes

(c) High-level protocols
- Transport Service
- Terminals
- File Transfer
- Job Transfer

A fourth Networkshop in this series has been arranged to be held at the University of York in April 1979. Although this event is also being organised by the Network Unit, by the time the meeting has taken place the Unit will have ceased formal existance, being replaced by the Joint Network Team. It is therefore appropriate to record an appreciation, on behalf of the university network community, for the expert and skillful work of the Director, Mervyn Williams, and of his colleagues Roland Rosner and Chris Morris. Successful networking requires cooperation at many levels between many bodies and individuals, and it is widely recognised that the Network Unit has played a vital and significant role in bringing about substantial cooperation in the university and research council community. We are indebted to them.

March 1979

D. F. Hartley
University of Cambridge

REPORT FROM THE NETWORK UNIT

Dr. R. A. Rosner

Cambridge Networkshop - September 1978

Report from the Network Unit

This paper is based on sections from the Network Unit's September report to the Computer Board and Research Councils. Where relevant, reference is made to projects and developments related to topics covered in that report.

The Adoption of Networking Standards

In April 1978, the Computer Board informed manufacturers that all machines delivered after 31 December 1980 would be required to conform with networking standards. (A copy of this operational requirement appears as an appendix to this report.) However, existing computers will also have to be connected to standard networks and the necessary adaption carried out without obligation on the manufacturers.

The standards of current concern are X25 Level 3, Triple-X (X3, X28 and X29) and the File Transfer Protocol (FTP). Pending a Post Office announcement on PSS, problems of interpretation or ambiguity in X25 will have to be resolved within the community. Such discrepancies as there may be between this version and that of the future PSS are unlikely to be serious. Subsequent changes for compatibility with PSS should therefore be minor. At the link level, HDLC chips are apparently still causing difficulties under some circumstances. There is, in any case, no urgency for the adoption of LAP or LAPB (at X25 Level 2) since other framing conventions, such as BSC, will be offered on PSS.

An informal group bringing together implementors of FTP is being run by Peter Linington (Cambridge). This discusses common problems and ensures that all implementations, including minimal versions, conform to the same detailed specification.

The academic community will undoubtedly require standards for several additional functions on a much shorter timescale than that envisaged for the output from the various standardisation bodies. It is important that, for these, all centres work to agreed definitions held centrally within the community.

The Connection of National Exporting Centres to PSS

The national exporting centres of the Computer Board and the Science Research Council will obviously be among the first to be connected to PSS. An informal group has accordingly been formed where these centres discuss common problems and attempt to coordinate their plans. These meetings have already indicated the need for definitions of conventions in certain areas. Since such decisions are likely to affect the whole community, it is important that there be interaction with other interested parties

and the Network Unit will organise this. A particularly important topic is the need for a sub-addressing scheme going beyond the limitations of the formal addresses in X25.

University Mainframes and PSS

Many university centres will wish their mainframes to be capable of connecting to PSS. The aim is, therefore, to build basic software/hardware packages for the most common combinations of machine ranges, operating systems and front-end processors thereby avoiding unnecessary duplication of work.

For a particular type of machine, there could be several possible approaches to the development of such a package, each involving different degrees of manpower and additional hardware. The handling of many networking functions is ideally carried out in close proximity to, if not actually within, the operating system. However, in a given case, the solution chosen will depend on the structure of the operating system, the severity of the modifications, the lifetime of the machine range integrated over the whole community and performance criteria. The Network Unit has recommended that the development of such packages should be centrally funded and carried out by placing contracts with centres in the community, manufacturers or software houses. Each project will consist of several steps:-

a) the presentation and evaluation of alternative approaches;

b) the choice of a suitable option followed by a detailed design study to produce a full specification;

c) the selection of a development site and a contractor together with estimates of manpower and additional hardware;

d) the development phase, with debugging and validation against centrally held standards;

e) the production of full documentation and maintenance arrangements with the contractor for the fixing of bugs and the introduction of improvements.

The installation of a package at a particular site must take into account the special nature of each configuration and its operating system and will make extensive demands on local expertise.

For some machine ranges, projects aimed at the production of networking packages are already in progress:

1) A proposal to develop an X25 connection for DECsystem 10 computers is being considered.

2) A joint ICL/universities Task Force has been established to consider networking requirements for 2900's.

3) A similar study for 1900's is about to begin.

Informal networking groups have grown around other machine ranges including GEC 4000's (under the auspices of Ken Heard at Harwell) and Prime. Pressure is growing for a group to be set up for PDP 11's and Daresbury have agreed to consider taking this on.

The integrity of each package is obviously of importance if it is to be widely distributed. Validation, documentation and maintenance are therefore crucial and some central control is needed to ensure that these functions are carried out to the required standard.

Science Research Council Networking

The SRC has two networks; one giving access to the IBM machines at Daresbury and Rutherford and the other for accessing the DECsystem 10's at UMIST and Edinburgh. The former network will be adapted to offer X25 ports by the end of the year and the IBM hosts will then be able to accept X25 calls. As a step towards merging the hitherto disjoint networks, the SRC is participating in the DEC 10 project mentioned above.

Daresbury and Rutherford belong to the informal group of National Exporting Centres.

Regional Networks

The lack of a public service and the absence of common standards have meant that the development of regional networks has necessitated a great deal of experimental work and heavy expenditure. Almost universally, the resulting products have proved to be highly dependent on proprietary hardware and software and, though capable of operating effectively in a fixed environment, have limited scope for evolution. They have also posed considerable control and managerial problems. Moreover, their ability to provide third-party switching puts them in contravention of the Post Office Act 1969. Though the PO has indicated a willingness to license existing arrangements, it is unlikely to sanction any new developments after the announcement of its own packet-switched service.

The Network Unit believes that the acceptance of CCITT standards and the expected availability of a public packet-switched service by 1980 provide a sound basis for inter-site communication in the future. It has, therefore, recommended that all prospective importers and exporters establish plans for connecting to the Post Office network. The resultant decoupling of communications from computing functions will offer a new flexibility, allowing

management and users to take decisions which are related to computing requirements rather than to rigidly defined communications links. This does not imply a total dependence on the PO service; changing traffic patterns and the pitch and structure of the packet-switched service tariffs may well point to economies through the use of private lines or regional switches rented from the PO. The use of standard protocols over such links will ensure general accessibility.

At the regional level, a distinction must be drawn between facilitating and resource-sharing networks. In the former, sites offer services on their own terms and the characteristics are determined by local management, users and hardware. In the latter, work is directed to specific machines in the interests of overall efficiency and there is a high commitment to the maintenance of service standards as perceived by the remote user. Although resorce sharing may make for a more rational use of the equipment in a region, the management overheads and the time spent in committee meetings to ensure the smooth functioning of such schemes can become intolerably high.

The Network Unit has recommended that future networks should be of the facilitating type and that the planned usage of remote resources should be based on bilateral agreements between importing and exporting sites.

Campus Networks

The computing environment on a university campus comprises a range of computing devices and terminal equipment and the objectives of a campus network are to allow connections to be established between pairs of entities on the campus and to provide a path to the outside world. Since it is at the campus level that the users and the communications arrangements provided by the various funding bodies come together, campus networks should play a key unifying role for computing in the university community.

In the future, the falling cost of hardware will lead to a greater diversity and distribution of computing resources. There will then be increasing demands from an expanding population of users for a variety of computing facilities and for significant improvements in the quality of services and convenience of utilisation. The availability of a campus network to which equipment and terminals can be connected simply, cheaply and speedily is a pre-requisite for achieving these objectives. Additional attributes are the capacity to carry large volumes of traffic (some of it flowing at high speed) and the possibility of using various protocols dependent on the nature of a particular transmission and the communicating entities. The Computer Board's Microprocessor Working Party considers that the effective exploitation of microprocessors will rely heavily on good

communications facilities. Software development for all levels of computer system is considerably simplified with the aid of a sophisticated filing system, a powerful editor and a lineprinter. It is not possible to provide such expensive facilities on every mini- and microcomputer system. Extensive cross-software development tools and debugging aids could be provided centrally on the campus for a range of machines and there could be shared access to expensive recording devices and other special peripherals.

The campus network must provide a gateway to the outside world both via the Post Office PSS and via direct links to major national centres where traffic volumes warrant them. Among the functions to be performed at the gateway are address mapping between local conventions and a nationally agreed scheme and the vetting of calls leaving the campus as protection against both unauthorised access to remote facilities and the generation of excessive traffic charges on the national network. The gateway's role will be simplified if due account is taken of national conventions in the design of the campus network.

There are several current projects to develop basic campus communications and switching systems some of which are described in Tony West's paper. The Network Unit has recommended that this work should be coordinated with the aim of creating standard commercially available products whose subsequent purchase should qualify for Computer Board funds.

Conclusions

Networking is an important aspect of Computer Board and Research Council computing policy. Over the past two years, there has been a growing acceptance of its potentialities and a gratifying awareness of the standards that are gradually emerging in this area.

It has been proposed that some form of central organisation be established to translate into action the ideas described in the foregoing sections. The basic premise is that progress is best achieved by cooperation among centres. The role of a central team would be to complement the activities in the various installations by tackling tasks such as fitting individual plans into a coordinated programme, supervising contracts for general packages, providing tools and facilities for implementors, disseminating information and liaising with the Post Office.

One of the problems affecting the evolution of networking is the extent of the demands on the expertise available in the community. This makes it all the more imperative that what is produced is of a very high quality and well documented so that installations can and will use each other's products.

Roland Rosner

Annex to Network Unit's Report to the September Workshop

Operational Requirements for a Computer in a Network Environment

Dr. R. A. Rosner

OPERATIONAL REQUIREMENTS FOR A COMPUTER IN A NETWORK ENVIRONMENT

Summary

Each machine will be required to communicate over a network with a diversity of computers and terminal equipment from a wide variety of manufacturers. The nature and number of devices participating in communications with the machine will vary dynamically from instant to instant thus demanding great flexibility in the operating system's mechanism for resource allocation.

Adherence to ratified national or international standards will be mandatory. Facilities must also be provided for users to supply their own modules to perform those functions for which standards are in the process of evolution.

Components

The order of the items in the list below corresponds to an architecture from the communications level to the job management level.

1. The machine will be connected to the network directly or via a communications controller or a front-end processor. Line speeds will be commensurate with the size and capacity of the machine. The control of these links must be in accordance with level 2 of X25 as specified in the technical guide for a future public packet-switched service provided by the Post Office.

2. The operating system must be capable of handling virtual calls and packets in conformity with level 3 of X25 as specified in the technical guide referred to above.

3. A module (or collection of modules) known as a transport station must be provided in the operating system to act as an interface between user processes and the network. The transport station will control the allocation of the network links as a multiplexed resource. It will be responsible for the implementation of a full end-to-end call mechanism and will offer an interface to user processes via an agreed set of process-to-process communications primitives. The transport station must conform to the specification currently being prepared by Working Group 3 of the BSI's Committee DPS/20.

4. The ability to communicate with terminals in accordance with CCITT Recommendations X3, X28 and X29 must be provided. Various high-level protocol handlers must be available particularly for the File Transfer Protocol published by the High Level Protocol Group working under the auspices of EPSS. All such protocols must be implemented as well-defined modules and there must be hooks to facilitate their easy replacement by other protocols either defined by the user or conforming to new standards as these emerge.

5. The manufacturer must indicate how his job control facilities may be exploited for the management of jobs in a network environment.

Associated with all the preceding components are accounting and statistics functions to monitor traffic and to journalise network and resource utilisation. Procedures must be supplied to carry out these measurements and to produce archival records.

Tools must be available for diagnostic purposes and for generating test environments to check out customer-written modules.

All components will be fully documented to allow easy usage.

Wherever possible, the manufacturer's internal protocols should be similar to those mentioned in 1 and 2 above.

Front-end processors must present no obstacles to the transparent transmission of bidirectional data streams so as to minimise the need for protocol conversion.

23 February 1978

An Introduction to Post Office Packet Switched Service

P. R. Morrison
Post Office Telecommunications

The information given by the speaker reflected the state at the time of the Workshop; in this section some documents reflecting the best information available at the time of going to press are reproduced. Details of the currently proposed tariffs and facilities are given.

The organisers wish to thank Pat Morrison for making this more recent information available to us.

PSS TARIFF SCHEDULE

There are 3 sections to the tariff structure:-

1 ACCESS CHARGES

2 USAGE CHARGES

3 FACILITY CHARGES

Where a connection charge is quoted this is a once only charge which is levied at the time that the item is provided. All charges in this schedule are VAT exclusive. VAT is payable on all Telecommunications Services and is included as an additional item on customer's bills.

1 ACCESS CHARGES

A packet mode terminal connects with PSS using a Dataline. This is a dedicated modem - line - modem - port connection operating at synchronous transmission speeds. There are four types of Dataline for packet mode terminals and the charge is independent of the relative locations of the terminal and the exchange.

TYPE	RENTAL £ P.A.	CONNECTION CHARGE £
Dataline 2400	1500	450
Dataline 4800	2500	700
Dataline 9600	3300	800
Dataline 48K	10000	2000

The minimum period of service for Dataline 48K is 3 years and for Dataline 2400, Dataline 4800, and Dataline 9600 is 1 year.

Character mode terminals connect to a PAD which is available at all exchange locations. Dedicated connection is available with two types of Dataline operating at asynchronous speeds.

TYPE	RENTAL £ P.A.	CONNECTION CHARGE £
Dataline 300	800	200
Dataline 1200	1100	350

Alternatively, character mode terminals may connect to the PAD using either Datel 200 (up to 300 bit/s) or Datel 600 (up to 1200/75 bit/s) and the Public Switched Telephone

Network. Normal charges for the exchange line, Datel Service Modem and the appropriate level of call charge to the PAD, will apply. The PAD locations are Birmingham, Bristol, Cambridge, Edinburgh, Glasgow, Leeds, London, Manchester and Reading.

In addition, character terminal users must input a Network User Identification as part of the input protocol. Each NUI is charged at £5 per quarter per NUI for each exchange at which the NUI is registered, subject to a maximum charge of £35 per quarter. It is possible for a user to have more than one NUI (for example, a means of identifying PSS usage by department) or several users may share one NUI (for example, a company may have many terminals at numerous locations and wish to operate a common NUI. There will therefore be one bill submitted by PSS against that one NUI). The connection charge is £25 per NUI.

PSTN access of PSS also involves the temporary occupation of a PAD port. This is charged at 65 pence per hour for Datel 200 and £1 per hour for Datel 600.

2 USAGE CHARGES

All PSS traffic is charged at one rate. There is no variation for distance or type of connection. All user data is carried in a packet structure. A packet consists of up to 128 octets and an octet may be considered as an 8 bit byte or 1 character. If the packet has 64 or less octets of user data (excluding the protocol overhead) then this is termed a segment and one unit of charge is made. Note that a packet with a zero data field is charged at 1 segment. If the packet has more than 64 octets of user data (up to the maximum of 128 octets) then 2 units of charge are made. Segments are totalled over a billing period i.e. 3 months, and the rate of charge is 23 pence per kilosegment (0.023 pence per segment or 0.00036 pence per octet).

PSS provides 3 types of interconnection between users:-

1 Virtual Call
2 Fast Select Virtual Call
3 Permanent Virtual Circuit

A more detailed explanation of the differences between these types of interconnection may be found in the facility schedule, but for tariff purposes it is necessary to state that the first two have a time element in the charge structure.

Both Virtual Calls and Fast Select Virtual Calls are timed for the duration of the call. All call times are totalled over the billing period and charged at the rate of 23 pence per hour. In addition there is a minimum usage charge made for each type of call; for Virtual Calls, all normal call attempts (other than those failing due to the Post Office network) will be charged 4 segments. On successful calls all customer data carried will be charged at the normal rate subject to a 4 segment minimum: that is, the total minimum charge for a successful call will be 8 segments.

For Fast Select Virtual Calls, all call attempts (other than those failing due to the Post Office network) will be charged 6 segments. On successful attempts all data carried in "data" packets will be charged at the normal rate.

The third type of inter-connection is a Permanent Virtual Circuit; this does not require a call set-up sequence and therefore there is no call duration element in the tariff. Instead there is a quarterly rental charge of £85 plus a connection charge of £4. The rental charge equates to virtual call duration of approximately 370 hours per quarter. Usage is charged at 23 pence per kilosegment.

3 FACILITY CHARGES

PSS aims to provide a wide range of facilities; a full description can be found in the facility schedule. The charges are listed below. In some cases a "change of status" charge is recorded, this is a once only charge made when the Post Office is notified of a change in facility status e.g. "please add this location to my Closed User Group".

If several changes are notified at one time then an abatement rule applies. This rule also applies to connection charges and reconfiguration charges. The rule states that £4 connection/reconfiguration/change-of-status charges will be reduced to £1 (other than the first) where several are ordered simultaneously for simultaneous execution.

Facility Charges for Packet Mode Users

1 Reversed Charge Acceptance: £4 per change of status.
2 Closed User Group: £4 per annum per CUG member terminal, plus £4 per reconfiguration.
3 Sub-Address: No charge.

4 Logical Channels: £4 per annum plus £4 change of status charge.

5 Multiple line transfer: separate dataline rental for each circuit.

6 Packet retransmission: retransmitted packets charged at the normal rate.

7 Redirection of calls: £5 per quarter plus £4 per change of address.

8 Window size selection: No charge.

9 Call statistics: 2 segments per call, plus £4 change of status.

Facility Charges for Character Mode Users

1 Reversed Charge Acceptance: £4 per change of status.

2 Closed User Group: £4 per annum per CUG member terminal per NUI plus £4 per reconfiguration.

3 Sub-address: No charge.

4 Network User Identification: £5 per quarter per NUI per exchange subject to £35 maximum. Connection charge £25 per NUI.

5 Direct Calling: £4 per annum plus normal call charges. £4 per change of address.

6 Editing: No charge.

7 User selection of PAD parameters: No charge.

8 Call statistics: No charge.

Note -

Fast Select Virtual Calls

A Packet Mode Terminal may initiate, receive and respond to Fast Select Virtual Calls if it subscribes to Fast Select Acceptance but a Character Mode Terminal may not initiate or respond to a Fast Select Virtual Call. A Dataline Character Terminal may receive Fast Select Virtual Calls if it subscribes to Fast Select Acceptance: there is a £4 charge for change of status.

01-432 9493 March 1979.

THE ANSWERS TO YOUR QUESTIONS ON PSS TARIFFS

Q1 What is the status of the tariffs in this paper?

A1 The PO has announced its plans to introduce PSS in March 1980. The tariff schedule associated with this paper is the tariff that will be levied for PSS.

No firm commitment can be made regarding the length of time that these charges will be held but every effort will be made to maintain these charges for the first year of operation.

Q2 What are the main elements of the tariff?

A2 For a packet mode synchronous access customer (X25) annual rentals are charged for access to the packet exchange and tariffs related to usage are charged for virtual call duration and data transmitted.

For character mode asynchronous customers (X3, 28 and 29) 2 types of charging are available:-

a. for Direct (Data line) connection an annual rental for access is charged, together with charges for virtual call duration and data transmitted;

b. for PSTN Dial-Up access: normal PSTN call charges and Datel Service modem rentals will be charged via current billing methods. Users will be charged by the hour for access to the network - "port holding charge" and also for virtual call duration and data transmitted. Dial-Up customers will be required to input a Network User Identification before data can be transmitted via the Packet Switched Service.

Q3 What are the charges for direct access?

A3 PSS will use the concept of a DATALINE at all speeds of access. The rental for a dataline includes a network port, 2 modems and the circuit from user to packet switching exchange. The rental for a dataline does not vary with the user's distance from the exchange.

Q4 What charge is made for port holding?

A4 This varies with speed and is:-

for 300 bit/s access 65p per hour
for 1200/75 bit/s access £1.00p per hour

This charge will be totalled over a quarter and rounded down to the nearest whole hour.

Q5 What is a NUI and is it a chargeable item?

A5 A NUI (Network User Identification) is the 'password' used to identify those customers accessing PSS via the PSTN and is chargeable.

A customer may rent several NUIs and may also share a NUI with other users.

NUI recognition may be:-

a. Limited to the home PSE

b. To any specified number of PSEs or

c. To all PSEs

Charges per quarter are for:-

a. £5

b.)
and) £5 per exchange up to a maximum of £35 per NUI
c.)

PSS bills will be sent to the renter of each NUI.

There is a connection charge of £25 per NUI.

Q6 What are the usage charges?

A6 These are in 2 parts: a virtual call duration charge, and a data transmitted charge.

The virtual call duration charge will be 23 pence per hour for each virtual call that a user initiates. The total hours will be summed over a quarter and the total, rounded down to the nearest whole number of hours charged.

A6 The data transmitted charge is calculated on the basis of the number of segments of 64 octets sent during a virtual call. A charge is made for each full or partially full segment and the charge of 23 pence applied to each complete thousand segments (kilosegment), summed over a quarter. An octet is 8 bits. Packets with zero octet data fields are charged as 1 segment.

These charges apply to all calls, whether initiated by a packet mode or a character mode terminal, and are distance independent.

Q7 Who pays these usage charges?

A7 Usage charges related to a particular virtual call are normally billed to the originator of the call, but can be billed, on a per call basis, to the other party if this is requested during call initiation and if the other party has indicated his willingness to accept. This is known as the Reverse charge facility and there is no extra charge for a Reverse charge call. But the Post Office does need to know if a terminal will accept Reverse charges.

Q8 Is there a minimum charge for a virtual call?

A8 All normal call attempts (other than those failing due to the PO network) will be charged 4 segments. On successful calls all customer data carried will be charged at the normal rate subject to a 4 segment minimum: ie the total minimum charge for a successful call will be 8 segments.

Q9 How about charges for other facilities?

A9 PSS will offer a large range of facilities. Some will be charged, others will not. For example, there will be annual charges for use of Closed User Groups and Logical Channels. These charges will not add more than 10% to a packet mode user's total bill; the average will be around 4%.

Q10 Do you offer Permanent Virtual Circuits (PVC's)?

A10 The PSS network offers the capacity to handle PVCs and the tariff that is proposed is a rental of £85 per quarter together with a kilo segment charge of 23 pence. The connection charge is £4.

Q11 Do you offer a Datagram Service?

A11 No. But we do offer an alternative. This is known as Fast Select and provides for user data in the call request and call clear packets.

All fast select call attempts (other than those failing due to the PO network) will be charged 6 segments. On successful attempts all data carried in 'data' packets will be charged at the normal rate.

Q12 What is the change of status charge and when is it applied?

A12 The Change of Status charge, where listed, refers to the charge raised when a customer either cancels his use of the facility or requests that the facility be made available to him. It does not apply when a customer is initially connected to the service.

Q13 How about 1200 bit/s duplex access?

A13 This is not currently available but we will announce tariff details once a PO modem for this speed of operation is available.

Q14 Will there be off-peak rates?

A14 We are continuing to examine the need for off-peak usage charges and port holding charge, but these do not at present form part of the tariff.

Q15 Will there be bulk discounts for high volume users?

A15 Not initially.

Q16 How will charges be billed?

A16 With the exception of PSTN users' modem and exchange line rentals and call charges for PSTN access to the network, all charges will be billed by a PSS billing centre. As for telephone service, rentals will be billed quarterly in advance; usage and port holding charges, quarterly in arrears.

A16 A bill will be raised for each PSS network address (for dataline access) and for each NUI (for PSTN Datel Service Dial-Up access).

Q17 Will national flat-rate PSTN access be available?

A17 No. Local and 'a' rate access to PSS will initially only be available to areas around the 9 PSS exchanges which are normally within those charge rates. PSS exchanges will be in London, Birmingham, Manchester, Bristol, Reading, Cambridge, Leeds, Glasgow and Edinburgh.

Users who would normally need to make 'b' rate STD calls to access PSS should note the availability of Dataline access to exchanges at 300 bit/s or 1200/75 bit/s.

The PO recognises the advantages of extending local rate PSTN access widely across the country and will do so as demand requires and facilities permit.

Q18 How does PSS compare with other PO services?

A18 PSS complements other PO data transmission services - the Datel Services - by providing higher speed switching and a wider range of facilities. PSS is more expensive than Datel for most simple point-to-point types of use but is likely to prove attractive for computer 'networking', distributed computing, terminal concentration and similar modes of operation. Potential users should note that some redesign of their system may be needed to fully utilize the facilities of PSS and to produce the maximum cost benefit.

Q19 How does this service compete with Dataplex, especially Dataplex 3?

A19 Both PSS and Dataplex 3 offer the remote character terminal concentration ability that is so attractive to timesharing bureaux. PSS offers easier extension of a bureau's geographical coverage, greater flexibility, less operational involvement - and hence cost - for the user and international access. The Post Office offers all its data services- side by side, and allows the Customer to make the choice.

NOTE: All charges listed in this document are VAT exclusive.
VAT is payable on all Telecommunications Services and is included as an additional item on customer bills.

PSS - THE PUBLIC DATA SERVICE

INTRODUCTION

This document aims to give an outline of PSS - the Service and its Facilities. It is intended for distribution to a wide cross-section of the public with varying degrees of understanding of packet switching. In view of this, a 2-tier approach has been adopted, with a simple explanation in the main paper and a more detailed description in the Appendices.

An explanation of starred terms is given in the Glossary at Appendix 1.

SERVICE DESCRIPTION

1 The service will offer 2 modes of operation viz:-

Packet Mode *
Character Mode *

Packet mode terminals * will connect to the service via a Dataline, a dedicated link which will comprise:-

The Modem in the customer's premises, the line, the Modem in the Packet Switching Exchange (PSE) and the access port.

Character terminals will connect to the service via either a Datel Service and the Public Switched Telephone Network (PSTN)* or a Dataline.

2 Datalines

2.1 For Packet Terminals there will be 4 types of Dataline:-

Dataline 2400
Dataline 4800
Dataline 9600
Dataline 48K

Each Dataline operates at the nominated synchronous transmission speed, eg Dataline 2400 operates synchronously at 2400 bit/s.

2.2 For Character Terminals there will be 2 types of Dataline:-

Dataline 300
Dataline 1200

Each Dataline operates at the nominated asynchronous transmission rate.

3 Character Terminal Access Using Datel Services

3.1 Character Terminals accessing PSS via the PSTN may do so by means of either the:-

Datel 200 Service
or
Datel 600 Service

3.2 Character Mode Terminals accessing the service via the PSTN will need to input a Network User Identifier * before calls can be made.

4 Terminal Interface

4.1 Packet Terminal Network Interface - Packet Terminals will have sufficient intelligence to interface to all three levels of the interface protocol.

4.2 Character Terminal Network Interface - Character Terminals will require assistance to communicate with other terminals on the network. This assistance is provided by PSE based equipment which will provide the necessary interface requirements as well as Packet assembly and disassembly.

Future reference to this equipment will be by its accepted abbreviation - PAD.*

5 Types of Call

5.1 The Network will support the following calls:-

5.1.1 Virtual Calls.*

5.1.2 Permanent Virtual Circuits.*

5.1.3 Fast Select.*

5.2 Packet Mode Terminals may, by means of either a Virtual Call, a Permanent Virtual Circuit or a Fast Select Call, communicate with other users as follows:-

5.2.1 Packet Terminal to Packet Terminal.

5.2.2 Packet Terminal to a Dataline Character Terminal.

Note - PSS offers Packet Terminals a full duplex data service, once a call is established data can be transmitted in either direction or both directions simultaneously.

5.3 A Character Terminal may communicate with other users by means of Virtual Calls only.

Character Terminal to Packet Terminal.

(From both Dataline connections and those accessing the network via the PSTN).

5.4 Call Restrictions

5.4.1 A Character Terminal may not call a Character Terminal.

5.4.2 A Packet Terminal cannot call a PSTN connected Character Terminal.

5.4.3 A Character Terminal may not initiate a Fast Select call or operate a Permanent Virtual Circuit.

6 Facilities

The following is a list of facilities available to PSS Users. For simplicity, they have been divided into those proper to Packet Terminals and those proper to Character Terminals. Detailed explanations are contained in the Appendices.

6.1 Packet Mode Facilities:-

i Reverse Charging and Reverse Charge Acceptance - An optional facility enabling the caller to request that the called party be charged for the call. Acceptance indicates the user's willingness to be offered such requests but the ability to reject any specific request is retained.

ii Closed User Group - Enables a group of users to inter-communicate but unless opting for one of the variants, precludes access by and to other network users.

iii Sub-Addressing - Allows the addition of 2 digits at the end of an address for call direction within a customer's terminal or communications complex.

iv One-Way Logical Channel* - Logical Channels may be assigned to operate in one of 4 modes as specified in Appendix 2.

v Multiple Line Transfer - Facility applies to the operation of multiple physical circuits between the customer and his PSE where the customer nominates a common address to all physical circuits.

vi Packet Retransmission - Allows a customer to request the retransmission of specified packets from the PSE to the customer's terminal.

vii Redirection of Calls - Allows, under certain conditions, for all incoming calls to be redirected to an alternative address if the called terminal is non-operational.

viii Window Size Selection - A packet flow control parameter.

ix Call Statistics - Subject to prior agreement, the network will provide details of call duration and data segment count at the end of a call.

6.2 Character Mode Facilities

i Reverse Charging and Reverse Charge Acceptance - An optional facility enabling the caller to request that the called party be charged for the call. Acceptance indicates the user's willingness to be offered such requests but the ability to reject any specific request is retained. Character Terminals using this facility are subject to the limitations explained in Appendix 3.

ii Closed User Group - Enables a group of users to inter-communicate but unless opting for one of the variants, precludes access by other network users. Operation of this facility by Character Terminals is subject to the network constraints for such terminals.

iii Sub-Addressing - Allows the addition of 2 digits at the end of an address for call direction within a customer's terminal or communications complex.

iv Calling Line Identification - Provides a Character Terminal with the identity of the calling customer.

v Direct Calling - Allows a Dataline Character Terminal to transfer data to one pre-arranged address without the need for call address information in the call set-up sequence.

vi Editing - Provides a very simple editing function of data held in the PAD prior to transmission of this date to the called customer.

vii User Selection of PAD Parameters - The configuring of the PAD to enable effective inter-communication between a Character Terminal and a PAD.

viii Call Statistics - All Character Terminals will automatically receive details of call duration and segment count for calls originated by them.

GLOSSARY

1 PSTN

Public Switched Telephone Network

2 PSE

Packet Switching Exchange

3 TERMINAL (DATA TERMINAL EQUIPMENT)

For the purpose of this document, a terminal may be:-

A host computer, communications controller, multiplexor, VDU, Teletype or other device which interfaces to PSS.

4 PACKET MODE TERMINAL

A data terminal equipment which can control and format packets, transmit and receive packets and has the intelligence to comply with Level 3 interface procedures. For PSS Packet Mode Terminals will operate synchronously at one of the following speeds:-

2400, 4800, 9600 or 48000 bit/s.

5 CHARACTER MODE TERMINAL

A data terminal equipment that does not have the intelligence to comply with Level 3 interface procedures and requires the assistance of a Packet Assembler/disassembler to communicate with other PSS terminals. It will operate asynchronously at - up to 300 bit/s or up to 1200 bit/s.

6 PACKET ASSEMBLER/DISASSEMBLER

Commonly referred to as a PAD. Performs the packet assembly and disassembly function for simple asynchronous Character Terminals. Also handles call set-up, clear and control procedures and generates service signals for transmission to the Character Terminal.

7 VIRTUAL CALL

A Virtual Call is the term applied to that period of time between call set-up and call clear, during which 2 users are effectively inter-connected for the transfer of data in either direction. There is no dedicated physical connection between the 2 parties, hence the term 'Virtual'.

8 LOGICAL CHANNEL

A Packet Mode Terminal may, subject to capabilities, support more than one Virtual Call, Permanent Virtual Circuit or Fast Select Call at a time. For the purpose of control such calls are transmitted or received over what are termed 'Logical Channels'. From the customer side of the line interface and similarly from the PSE side of the line such channels are seen as individual channels and are identifiable within a contiguously numbered range. However, data from each logical channel is transmitted across the line in sequence but interleaved with data from other channels. A Character Terminal can only support one Logical Channel.

9 PERMANENT VIRTUAL CIRCUIT

A Packet Mode customer may, in agreement with the network, elect that all data transmitted across a designated Logical Channel should be directed to a pre-determined address. Such a channel is always in the data-transfer state and hence it is not necessary to establish a Virtual Call before transmitting data.

Logical channels so assigned are termed 'Permanent Virtual Circuits'.

10 OCTET

Octet is the term applied to a group of 8 bits of data.

11 FAST SELECT AND FAST SELECT ACCEPTANCE

A Fast Select call provides for the inclusion of up to 128 Octets of user data within a call request sequence and similarly in the response from the called customer. Fast Select Acceptance indicates a willingness to receive Fast Select Calls.

In outline the facility operates in the following mode:-

> The originator sends a Call Request packet to the network and includes with it up to 128 Octets of 'User data'. The recipient may then reply in one of 2 ways. He either generates a Call Clear packet to which he may, if he wishes, add a similar 128 Octets of user data. This would then terminate the call. However, should the recipient wish to extend the call, he may respond with a Call Accept packet plus of course 128 Octets of data, if he wishes. The effect of this response is to convert the Fast Select call into a Virtual Call for which normal charges will apply.

Character Terminals may not initiate Fast Select calls but a Dataline Character Terminal may agree to receive such calls. It should be noted however, that the Dataline Character Terminal may not respond to the incoming Fast Select Call which will be cleared automatically, by the network, after delivery of the user data.

12 EXTENDED FORMATS

Although not a specific facility, the use of Extended Formats is required in conjunction with other facilities such as Window Size Selection, Re-direction of Calls and Call Statistics.

13 NETWORK USER IDENTIFICATION (NUI)

Character Mode Terminals accessing the service via the PSTN will need to identify themselves before the network will allow calls to be made. The identifier will take the form of 12 alpha-numeric digits of which the first 6 will be chosen by the customer and the second 6, will be allotted by the Post Office. The last 6 digits will on input to the PSE be rendered illegible on the customer terminal if it is working in full duplex or is a half-duplex terminal with a back-space capability.

A customer may elect to use one or more NUIs and for each of these to be recognised at one, any or all PSEs.

FACILITIES - PACKET MODE TERMINALS

1 REVERSE CHARGING

1.1 Reverse Charge Request - This is an optional User facility available on a per call basis. It enables a user, when setting-up a call, to request that the usage charges ie data transmitted and duration charges, for that call, be debited to the called party. However, such calls will only be connected if the called party has previously indicated his willingness to accept such calls.

1.2 Reverse Charge Acceptance - A customer is required, at subscription, to indicate whether he will accept incoming call requests for reverse charging. As indicated previously, the called party has the option on a per call basis, to reject such call requests by generating a call clear sequence.

In the absence of an agreement to the acceptance of calls requesting reverse charging, the network will initiate a call clear sequence to the calling customer.

1.3 Customers will not be able to either request or receive International reverse charge calls.

2 CLOSED USER GROUP (CUG)

This facility in its basic form enables a group of users to communicate with each other across the network whilst at the same time precluding communication to or from other subscribers on the service. However, for some users the basic form may be over-restrictive and individual subscribers may belong to up to 99 CUGs; each group being identified by a separate 2 digit numerical code. Alternatively users may opt for one of the variations listed in Paras 2.2, 2.3 or 2.4 below:-

2.1 Basic Form - As described above.

2.2 Closed User Group with Outgoing Access - Individual members of a CUG may, if they wish, also have access to either other Closed User Groups or to other Users in general, subject to any limitations of the service.

2.3 Closed User Group with Incoming Access - Individual members of a CUG may, if they choose, receive incoming calls from either other CUGs or from other network users.

2.4 Closed User Group with Outgoing and Incoming Access - This variation combines the facilities described in Paras 2.2 and 2.3 above.

3 SUB-ADDRESSING

It is anticipated that some customer terminals may form part of a communications complex and as such could present the need to further identify the address. Provision has therefore been made for the addition of 2 sub-address digits to be added immediately following the called address. This address information will be carried transparently across the network between the calling and called parties and therefore any response to this information rests entirely with the recipient.

4 ONE-WAY LOGICAL CHANNEL

Logical Channels may be assigned for usage in one of 4 modes. Each mode constitutes a 'Logical Channel Group' and each 'Group' is allocated a contiguous number range. There are both theoretical and practical limits to the total number of channels that can be operated within each Group; the theoretical limit is unlikely to be either required or offered and the practical limit has yet to be set.

The 4 modes are as set out below and whilst the order relates to allocation in practice, not all types need to be utilised:-

4.1 Permanent Virtual Circuit - A full description of this facility has been give in Appendix 1.

4.2 Incoming Only - Incoming calls will be allocated to the lowest numbered free channel. If all channels are busy then calls will be offered on a 'Bothway Channel'.

4.3 Bothway - Again, incoming calls will be offered on the lowest numbered free channel.

4.4 Outgoing Only - The customer's terminal may elect to use either an Outgoing or a Bothway channel as it requires, when both are available.

5 MULTIPLE LINE TRANSFER

A customer may elect to have a number of independent physical circuits between his terminal and the PSE but at the same time allocate to them a common network address.

When PSS Service opens, lines will operate with individual Level 2 (link access) and Level 3 (packet level) procedures. However, it is expected at some future date, to be able to operate a common Level 2 and Level 3 procedure across such multiple physical circuits. At that time, the network will be able to support either mode of working.

6 PACKET RE-TRANSMISSION

This is a user facility that operates, at Level 3, between the customer's terminal and his PSE and should not be confused with the control procedures that operate within the network. It enables the customer to request, within certain limits, the re-transmission of specified data packets. Re-transmission of the specified packets must be complete before further requests for re-transmission can be sent. Failure to observe this procedure will result in the call being reset which means that although the call is not cleared, short-term stores are cleared and thus the packets for which re-transmission is sought, will be irretrievable.

7 RE-DIRECTION OF CALLS

A customer may register, with the network, an alternative address to which all incoming calls should be directed if and when his terminal is non-operational. If the call is proper to a member of a Closed User Group, then re-direction may only be to another member of the same Closed User Group.

The network will not attempt to re-direct a call more than twice.

8 OTHER FACILITIES

The following facility although forming part of the list of facilities available to Packet Mode Terminals, requires an in-depth understanding of packet switching and therefore has not been described in detail:-

8.1 Window Size Selection - A packet flow control parameter which, in the absence of a customer setting will be allocated a value by the network.

9 CALL STATISTICS

Customers may if they choose, elect to be advised of the number of data segments transferred and the duration of each call. This information will be sent to the originator by the network at the conclusion of each call, regardless of whether the call is cleared by the calling or called party.

FACILITIES - CHARACTER MODE TERMINALS

1 REVERSE CHARGING

1.1 Reverse Charge Request - This is an optional User facility available on a per call basis. It enables a User, when setting up a call, to request that the usage charges ie data transmitted and duration charges, for that call, be debited to the called party. However, such calls will only be connected if the called party has previously indicated his willingness to accept such calls.

1.2 Reverse Charge Acceptance - For Character Terminals, this facility is restricted to those terminals connected to the service by Datalines and must conform to the following requirements.

The customer is required, at subscription, to indicate his willingness to accept incoming calls for which he may be liable for the charges. However, this does not preclude the ability to terminate any call which he does not wish to accept.

In the absence of an agreement to accept such call requests, the network will generate a call clear indication to the calling customer.

1.3 The restriction on the use of Reverse Charge Request and Reverse Charge Acceptance in conjunction with International calls, also applies to Character Terminals.

2 CLOSED USER GROUP (CUG)

This facility in its basic form enables a group of users to communicate with each other across the network whilst at the same time precluding communication to or from other subscribers on the service. However, for some users the basic form may be over-restrictive and individual subscribers may belong to up to 99 CUGs; each group being identified by a separate 2 digit numercial code. Alternatively, Character Terminals may opt for one of the following variants:-

2.1 Character Terminals accessing the network via the PSTN will be limited to operate in the basic mode and permitted then only to generate calls to Packet Mode Terminals within the Group, or to have additionally outgoing access to other users.

2.2 Character Terminals connected to the service by a Dataline may select from the four variations of Closed User Groups (as listed in para 2 Appendix 2) and are limited only by the restriction that they cannot communicate with other Character Terminals.

3 SUB-ADDRESSING

It is anticipated that some customer terminals may form part of a communications complex and as such could present the need to further identify the address. Provision has therefore been made for the addition of 2 sub-address digits to be added immediately following the called address. This sub-address information will be carried transparently across the network between the calling and called parties and therefore any response to this information rests entirely with the recipient.

The operation of this facility by Character Terminals is subject only to the limitation that Character Terminals may not inter-communicate.

4 CALLING LINE IDENTIFICATION

This facility is only available to Character Terminals connected to the network by a Dataline and identifies, by means of the National PSS number, the originator of an incoming call.

5 DIRECT CALLING

Character Terminals connected to the service by a Dataline may, by prior agreement, request establishment of a call to a pre-arranged network address without the need to specify the called customer's address. The facility is roughly equivalent to a permanent virtual circuit but is restricted to a single agreed address.

The provision of this facility does not preclude the use of the terminal for calls to other destinations and for this the addressing information must be included in the call request sequence.

6 EDITING

The PAD will provide very simple editing functions of the data held in its short-term store, prior to transmission to the called customer. The editing function will be:-

6.1 Deletion of the last character.

6.2 Deletion of the last 'message'.

6.3 Display of the last 'message'.

The word 'message' should not be taken literally. In this context it may merely be the last line or 2 lines. Further explanation of the functions covered in Para 6.2 and 6.3 would necessitate in-depth understanding of packet switching techniques and therefore detailed explanations have been omitted.

7 USER SELECTION OF PAD PARAMETERS

This is perhaps one of the most complex aspects of PSS and therefore no attempt will be made to give more than an outline of the facility in this document.

For a Character Terminal to communicate effectively with the PAD and vice versa, it is necessary to define the conditions and responses one can expect from the other. Terminal characteristics are of a fixed nature and therefore the PAD adapts to these to enable inter-communication to take place.

In practice each Character Terminal type is identified by a 2 digit code which, when used during the log-in sequence of a call, via the PSTN, allows the PAD to re-configure its parameters and so communicate with the calling party.

8 CALL STATISTICS

At the end of each call originated by a Character Mode Terminal, the network will send, to the call originator, details of the duration and segment count for that call. This facility will be provided automatically by the network and will incur no charge.

Permission to Connect (PTC)

Technical Standards for Procedural Interfaces

K. Knightson
Post Office Telecommunications

STOCKHOLM, 13-21 SEPTEMBER 1978

SOURCE : UNITED KINGDOM POST OFFICE

TITLE : PERMISSION TO CONNECT (PTC)
TECHNICAL STANDARDS FOR PROCEDURAL INTERFACES

1 INTRODUCTION

This contribution is concerned with the requirements that should be applied to the procedural aspects of DTE/DCE interfaces during permission to connect (PTC) exercises.

The criteria for safety and electrical compatability are already under discussion within CD/SE in the context of IEC 435 and are thus not covered in this contribution.

2 BACKGROUND

The UKPO's Experimental Packet Switched Service (which first began operation in late 1975) has provided the UK Post Office with considerable experience in the problems that may be encountered in dealing with the procedural aspects of DTE/DCE interfaces.

One important feature of the DTE/DCE interface that became apparent at a very early stage was that because of the nature of the protocols neither the DTE nor the DCE can fully operate without being connected together. This clearly makes the task of locating any procedural problem very difficult. Furthermore, the nature of computer controlled equipment is such that it cannot always be categorically claimed that the fault is the customers DTE, and never with the PTT equipment.

The nature of the protocols is also such that neither the DTE or DCE can be successfully looped to themselves for self testing.

Another problem that became apparent was that the speed and combination of events can become so complex that they cannot be tracked directly by a human. Recording cannot overcome this problem since examination of long recordings is an extremely tedious process and prone to error.

Direct examination (line by line) of the customers software is clearly not a viable solution either for reasons similar to those above or the multitude of programming languages and further interactions with other system software.

3 THE NECESSITY FOR 'APPROVAL' OF PROCEDURES

If a DTE meets the safety and electrical compatability criteria the PTTs must decide whether they also wish to examine in detail the procedural operation of the DTE, bearing in mind the resources that will be required, and the corresponding benefits to be gained.

From EPSS the UK has found that the inability to adequately test the DTE/DCE interface procedures leads to:-

a. Degradation of network performance as seen by other users.

b. Potential maintenance liabilities

c. Unresolvable demarcation disputes.

In most cases it may be that the causes of a.-c. lie within the DTE. However, the inability of the PTT to assist in resolving these issues reflects badly on the PTT irrespective of where the fault lies, and thus eventually the customer loses confidence in the use of public switched services, since he gets no service irrespective of whether it is his supplier's fault or the PTT's fault.

At a time when great emphasis is being placed within ISO on 'Open System Interconnection' (which is a drive for universal compatibility of a wide variety of high level protocols) the PTT must do all they can to assist the user community to make use of public switched services.

4 PROCEDURAL TESTING

Most procedures (sometimes called protocols) can be divided into layers or levels. For example, X25 comprises 3 major distinct levels, 1, 2 and 3. Furthermore each individual major level can usually be sub-divided into further sub-levels. For example, level 2 of X25 can be divided into the frame structure level and the elements of procedures. The elements of procedures can again be further sub-divided into phases of link-set-up, information transfer, disconnection etc.

The PTC exercise could be simplified by using this technique of dividing the procedures into definable separate parts. The same divisions of course would be used for subsequent maintenance and demarcation problems. It is the firm belief of the UK that this approach is vital in avoiding the problems that complex procedures create and that complete chaos will occur unless this approach is chosen.

5 RECOMMENDATIONS

The UK recommends that urgent effort be expended on dividing up procedural interfaces into small defined levels and sub-modules, small enough to enable testing and event monitoring to be achieved at the speed of human operation.

Within each defined smallest module the DTE must be able to initiate a single event or series of events on demand by the testing officer, and be able to respond with single or multiple response again on demand by the testing officer subsequent to a PTT generated event. Adequate manual input and output facilities must be available. A range of testing configurations is foreseen as being necessary, among which could be included

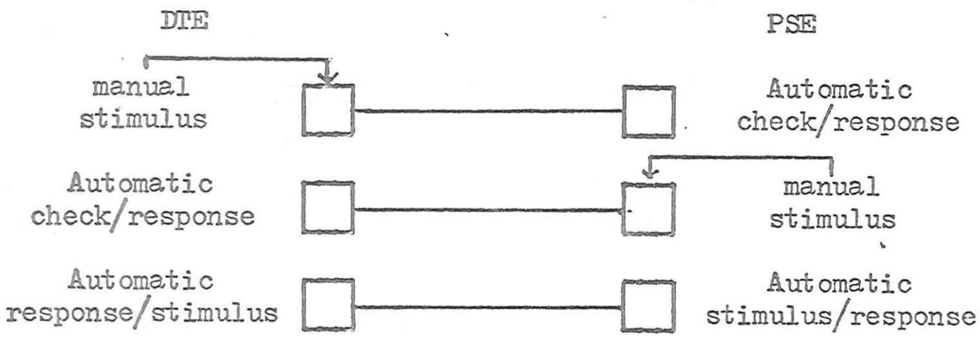

These principles can be applied to any given procedural interface. The UK has so far considered X25 and has in mind the following types of tests.

5.1 LEVEL 2

5.1.1 Generation by the DTE of any given frame type on demand.

5.1.2 The reporting of receipt of any frame type and its component values.

5.1.3 A complete link set-up sequence test.

5.1.4 Generation of sequences of I frames with required component values (possibly deliberate violation).

5.1.5 Reporting of received sequences of I frames with component values (possibility deliberately violated by the testing officer).

5.1.6 Tests 5.1.4 and 5.1.5 run simultaneously for appropriate action/reaction (automatic).

5.2 LEVEL 3 (assuming level 2 now tested and connected)

5.2.1 Generation of any required level 3 packet.

5.2.2 Reporting reception of any level 3 type packet.

5.2.3 Generation of data packets with required sequence numbers with or without deliberate violations.

5.2.4 Reporting of received data packets and sequence number reporting (including induced violations).

5.2.5 Tests 5.2.3 and 5.2.4 run together to test for appropriate action/reactions (automatically).

5.2.6 Facility field tests, address tests, M and Q tests, cause fields, channel numbering etc.

6 IMPLICATIONS FOR THE PTT

In order to perform tests of the nature recommended the PTTs will require sophisticated test equipment and manpower resources. The UK intends to reduce the manpower resources by having the test equipment located at the exchange sites, rather than taking the equipment to the customers premises. This has an additional advantage that several customers can be tested simultaneously, and that if the tests are properly defined then the customer can develop and test his equipment himself using the PTT provided test-bed prior to formally requesting the PTT to perform the PTC exercise. If the test-bed itself keeps records then the PTT will automatically know when the customer equipment is satisfactory. Clearly the implementation of this facility may vary from PTT to PTT but this section gives an outline of the possibilities.

The UK intends to co-locate the test-bed with the exchange and then divert the customer lines into the exchange when the DTE has been given PTC. Lines could be re-diverted back to the test-bed for subsequent maintenance and demarcation disputes etc. In addition it is intended that the test-bed can also test the exchange and thus arbitrate in demarcation disputes. The UK estimates that this

approach will reduce the man-power resources expended visiting customers' premises and will in fact be cheaper in the longer term.

7 IMPLICATIONS FOR THE DTE MANUFACTURER/IMPLEMENTER

Enforcement of the tests outlined above clearly creates extra work and overheads for the DTE manufacturer/implementer. It requires the inclusion, albeit on a roll-in/roll-out basis, of special routines that would not be part of the normal operational environment.

However, many of the UK manufacturers/implementers have already expressed interest in the provision of such test-beds because it can expedite their own development and greatly assist in debugging their DTEs.

8 SUMMARY

The UK recommends that procedural interfaces should be included in the PTC exercise and should be divided into small enough defined elements for ease of testing on a manual basis. Even though this requires extra resources on the part of both the PTT and the DTE manufacturer/implementer this is more than offset by the benefits gained.

Finally, to cater for cases where designs already exist and cannot (or the manufacturer is unwilling to) include the required test facilities, further study is necessary to determine whether a common CEPT attitude to such situations is possible or desirable. Alternatively, individual PTTs may need to operate their own policies in relation to whether or not connexion of such DTEs would be allowed, and this might involve, for example, the withholding of PTT assistance in cases where the DTE encounters procedural difficulties.

Clearly, agreement to the UK approach within CEPT would provide manufacturers with a harmonised approach and make their equipment maintenable throughout the CEPT countries and proportionately reduce the overheads incurred in providing the extra facilities required.

Network Topology 1981

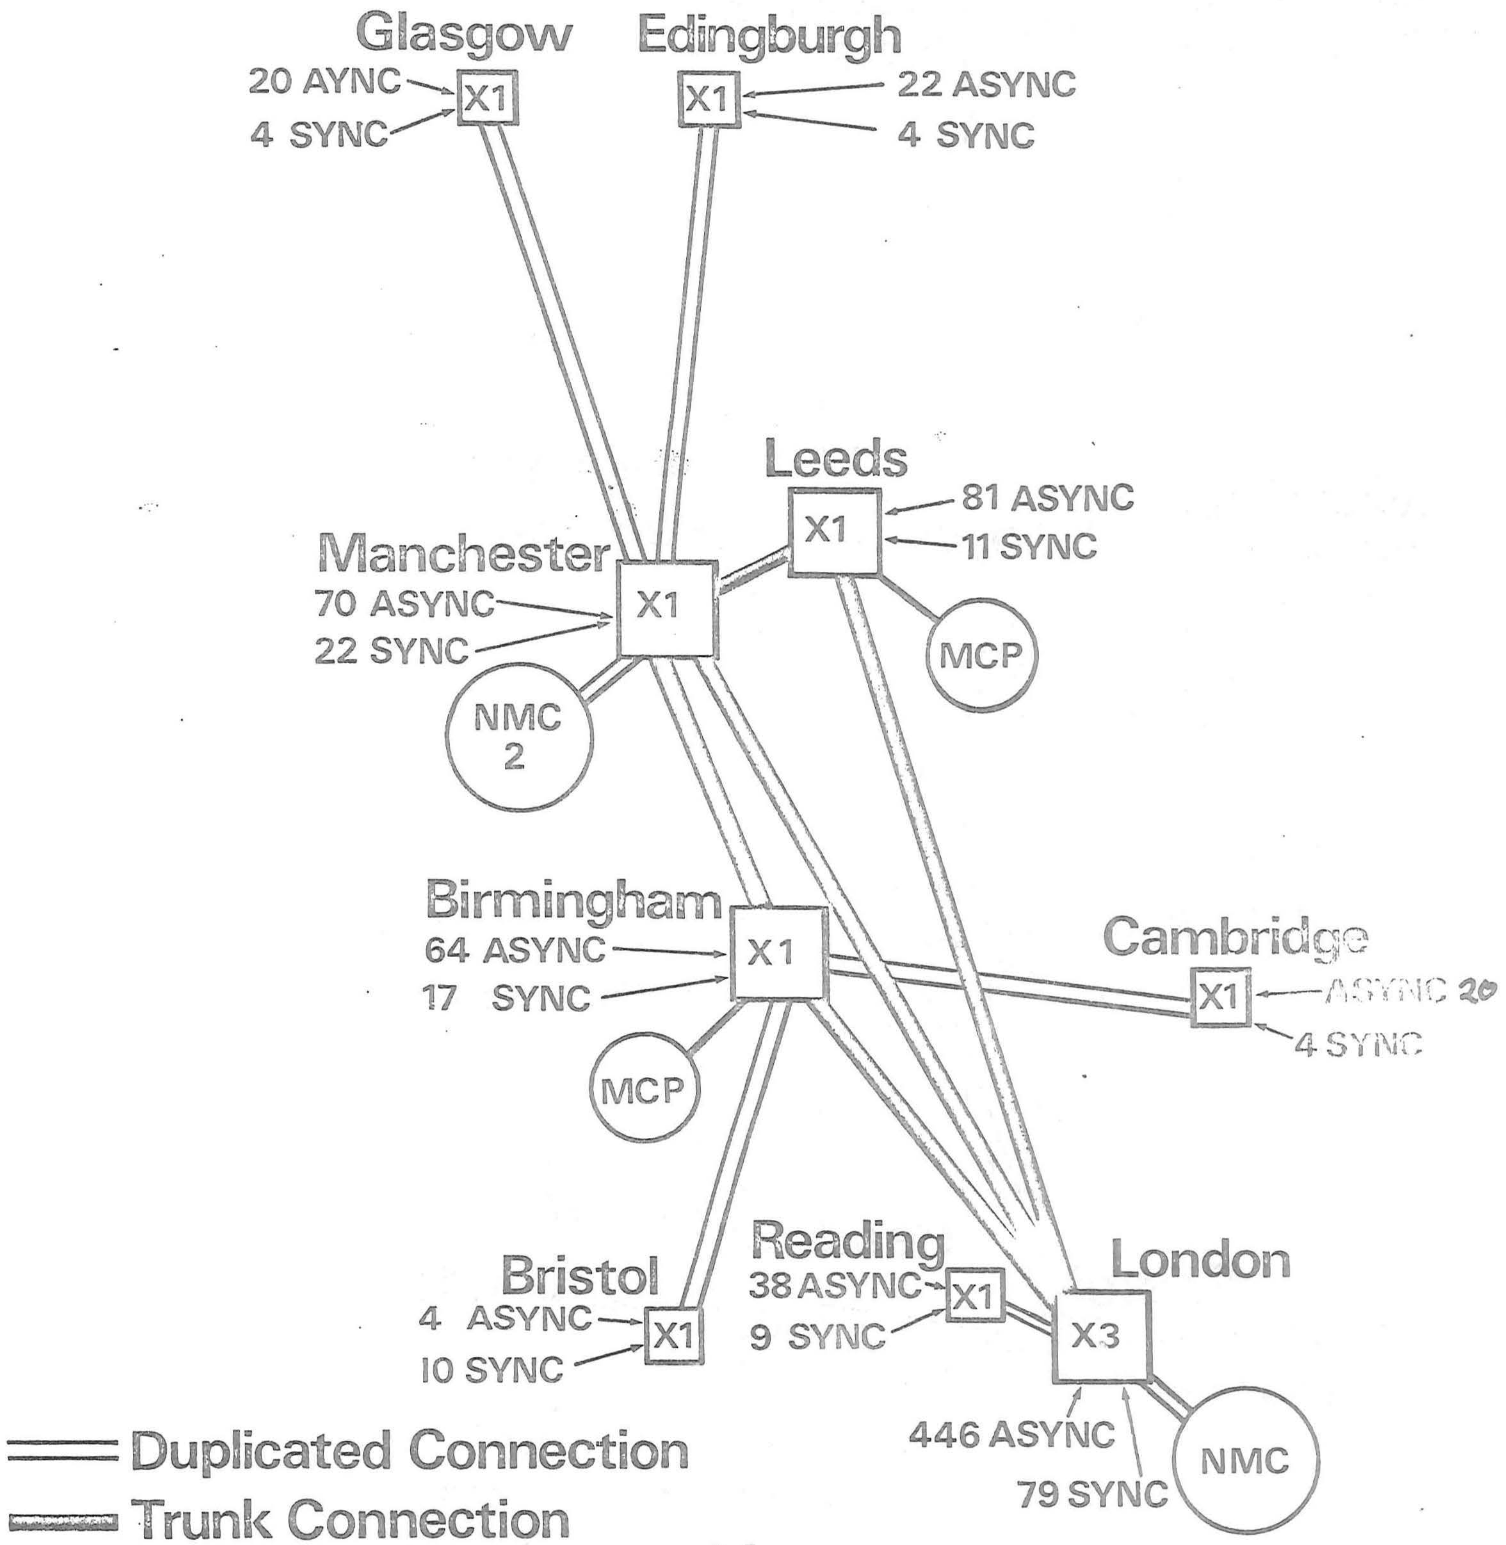

Network Topology 1983

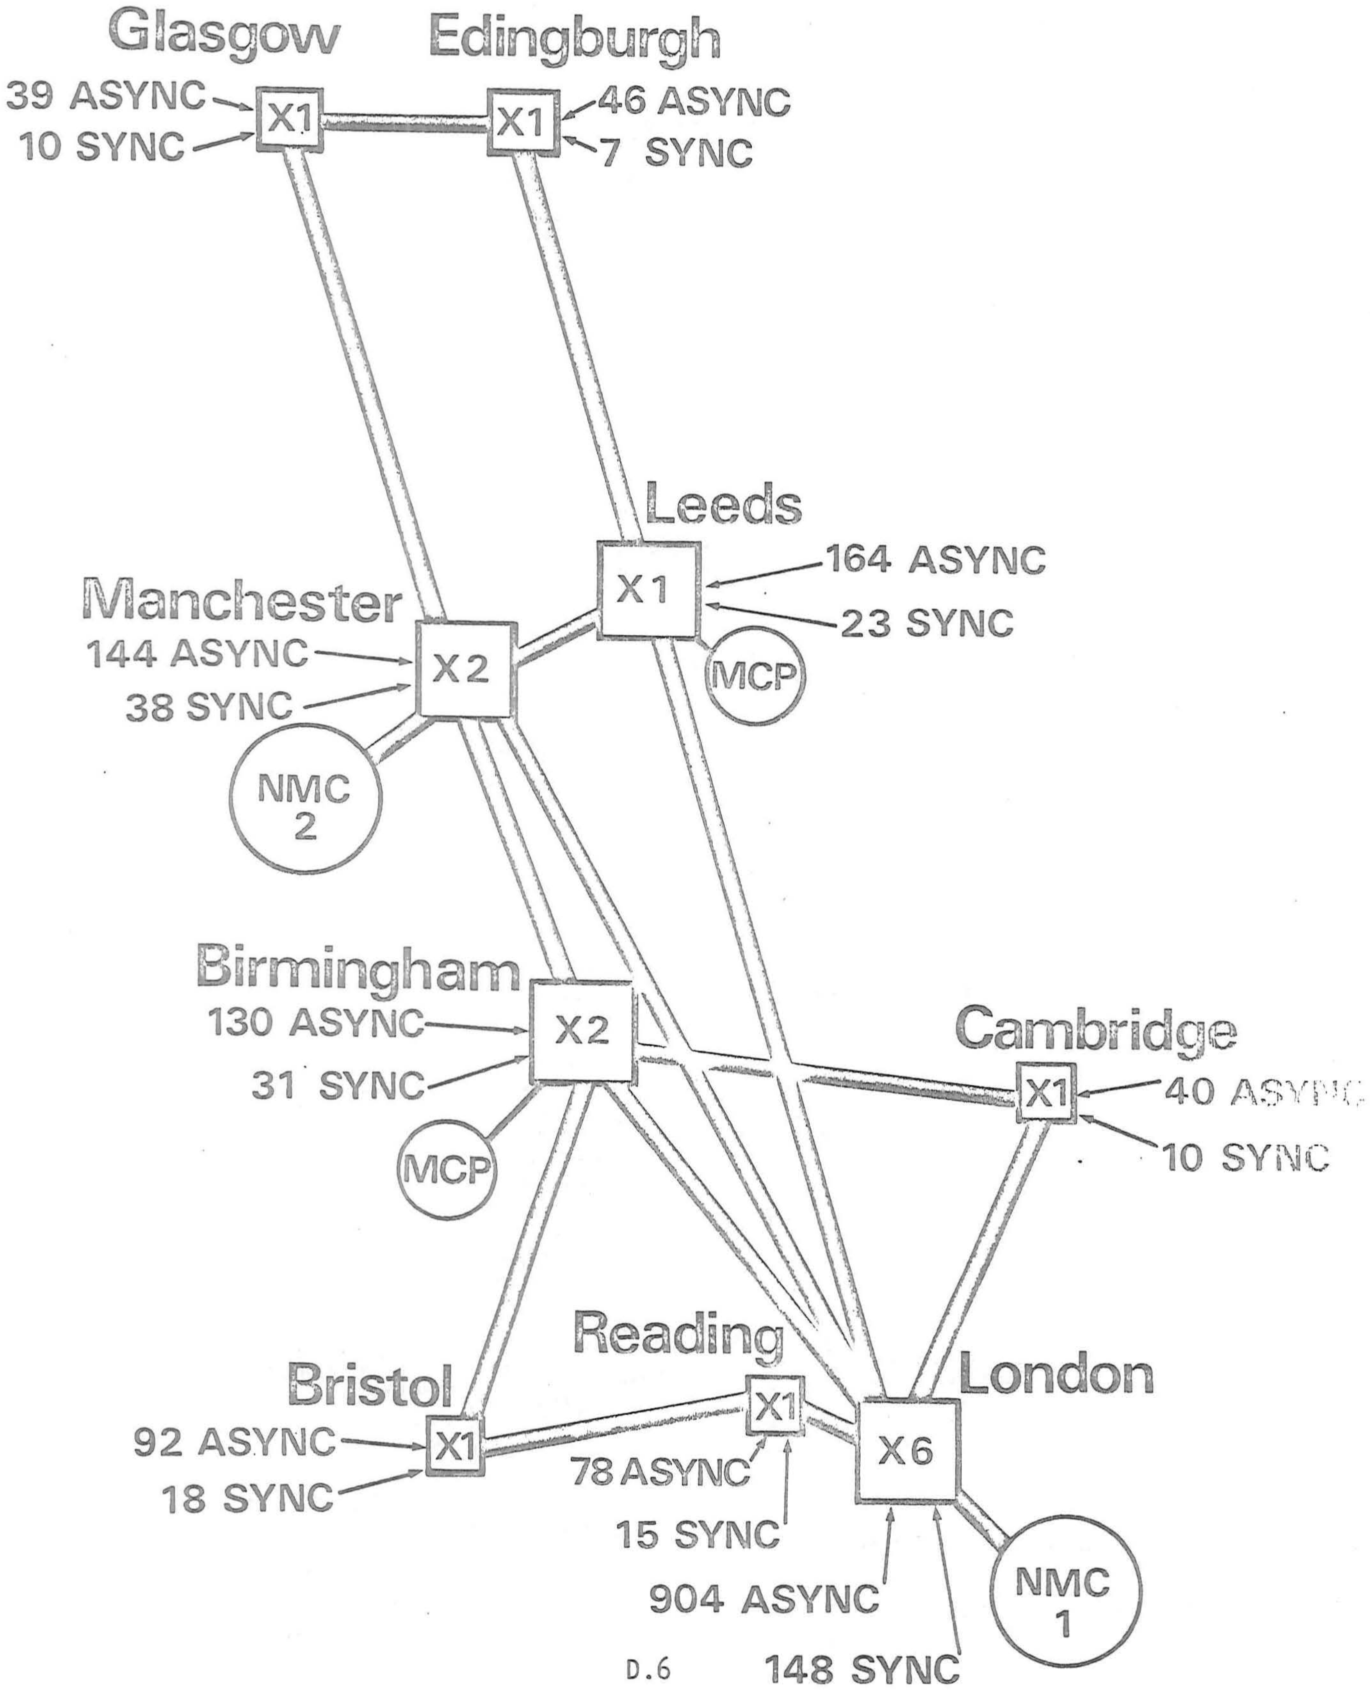

PSS

X25 LEVEL 2

1. LAP

Essential Symmetric

2. LAPB
3. BSC
4. Multi-line Transfer

PSS

X25 LEVEL 3 ESSENTIAL

1. Reverse Charging
2. Reverse Charge Acceptance
3. Closed User Groups
4. One-way Logical Channel
5. Sub-addressing (2 Digits)
6. Packet Retransmission
7. Multiple Lines With Single Address

PSS
X25 LEVEL 3 POSSIBLE

1. Packet Size Selection

2. Window Size Selection

3. Fast Select

4. Multi-line Transfer

5. Priority Selection

PSS
X28 FACILITIES ESSENTIAL

1. NUI (IPSS & EURONET Compatible)

2. Reverse Charging

3. Reverse Acceptance (Directly Connected)

4. Closed User Groups

5. Sub-addressing

6. Calling Line Identification

PSS

X28 POSSIBLE FACILITIES

1. Direct Calling
2. Editing
3. Multiple Asychronous Host Connection
4. Extra X3 Parameters

TESTING AID PHILOSOPHY
PSE SITE

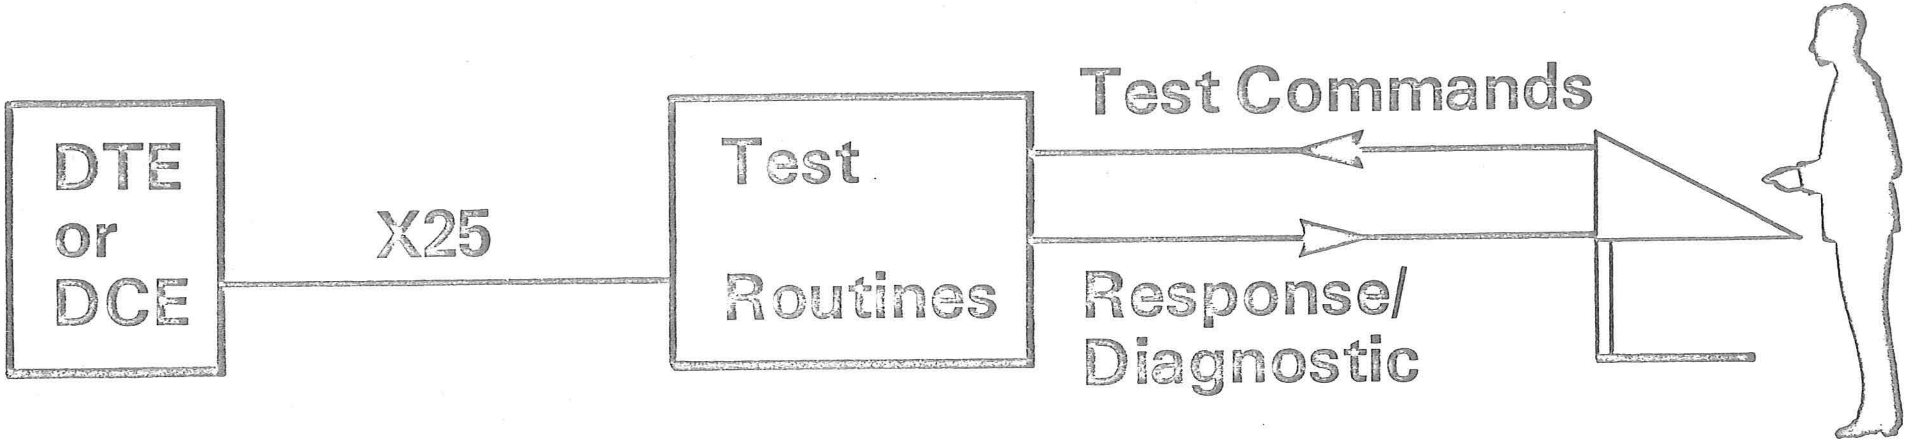

1. DTE Development/Debugging

2. Demarcation Resolution

3. Permission To Connect (Approval)

Interim Testing Aids

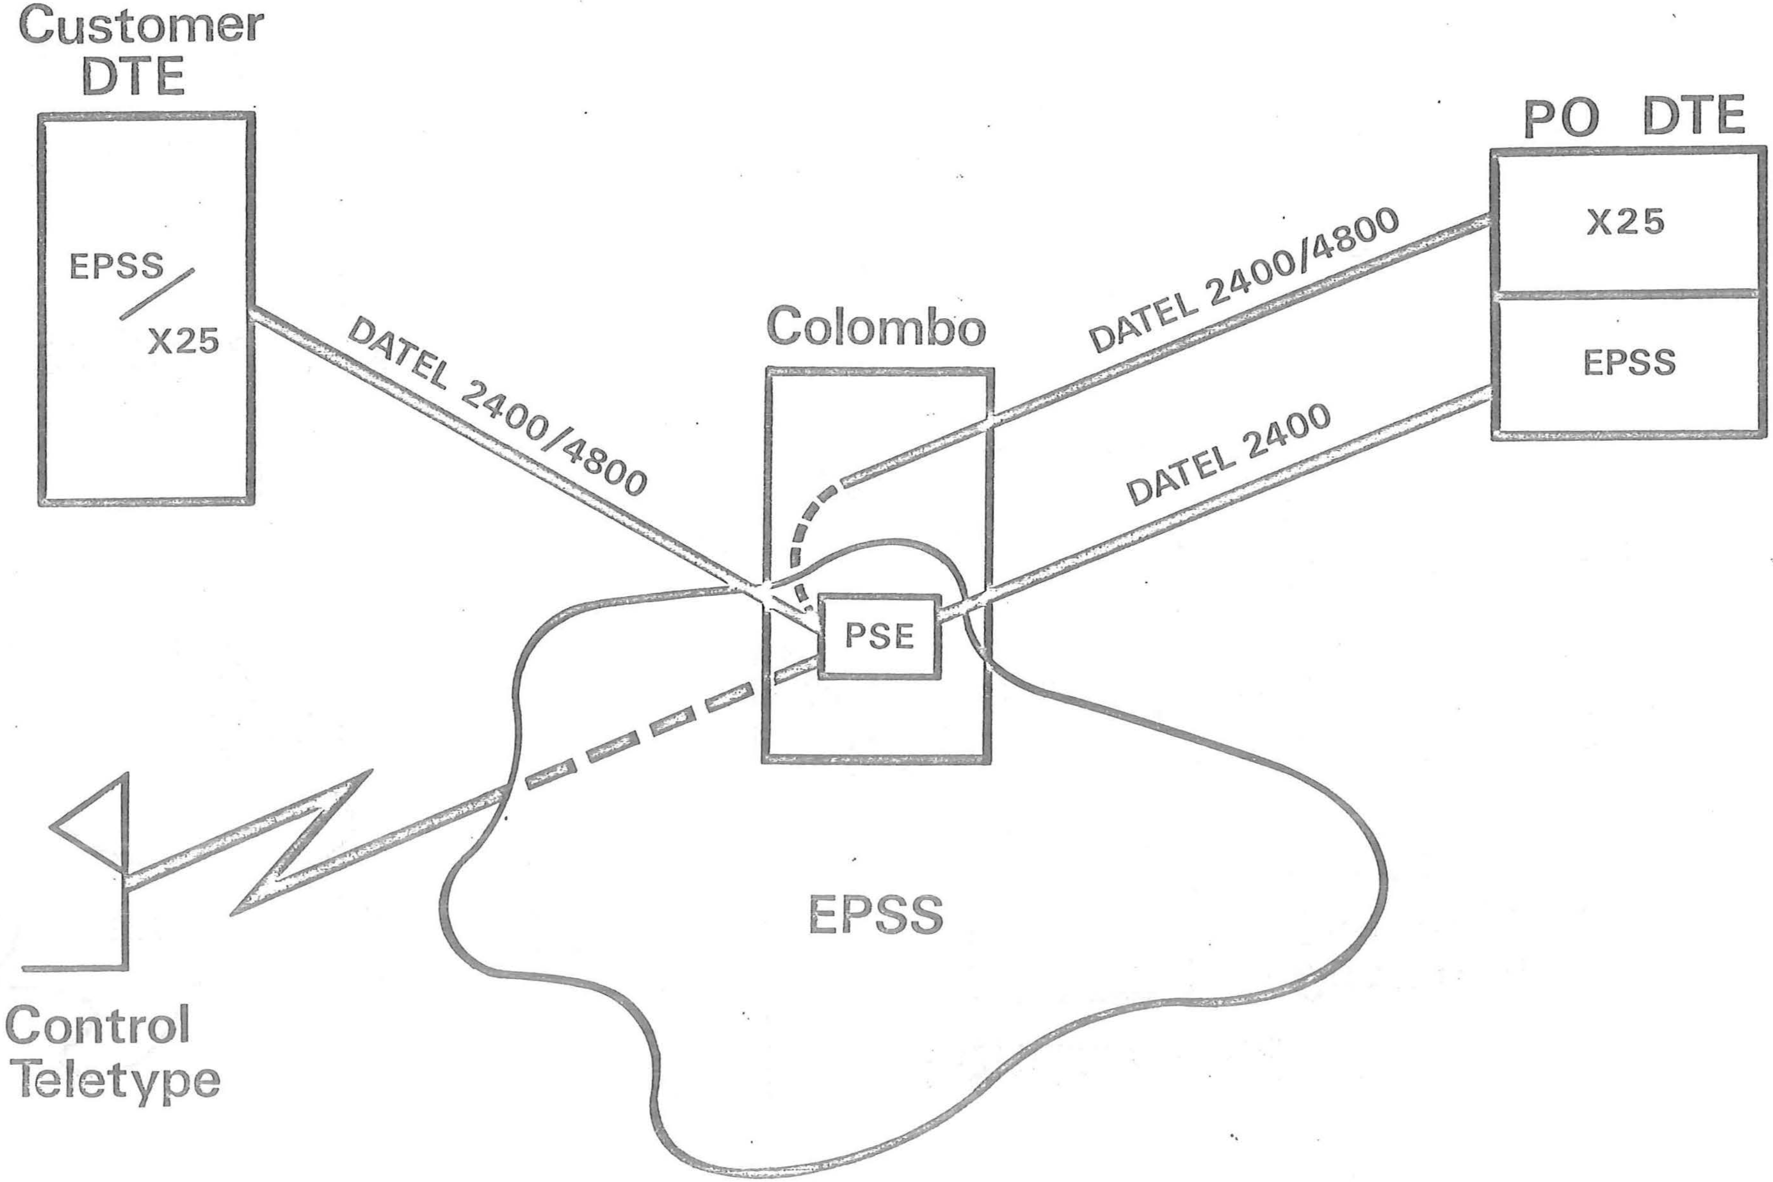

Medium Term Testing Aids

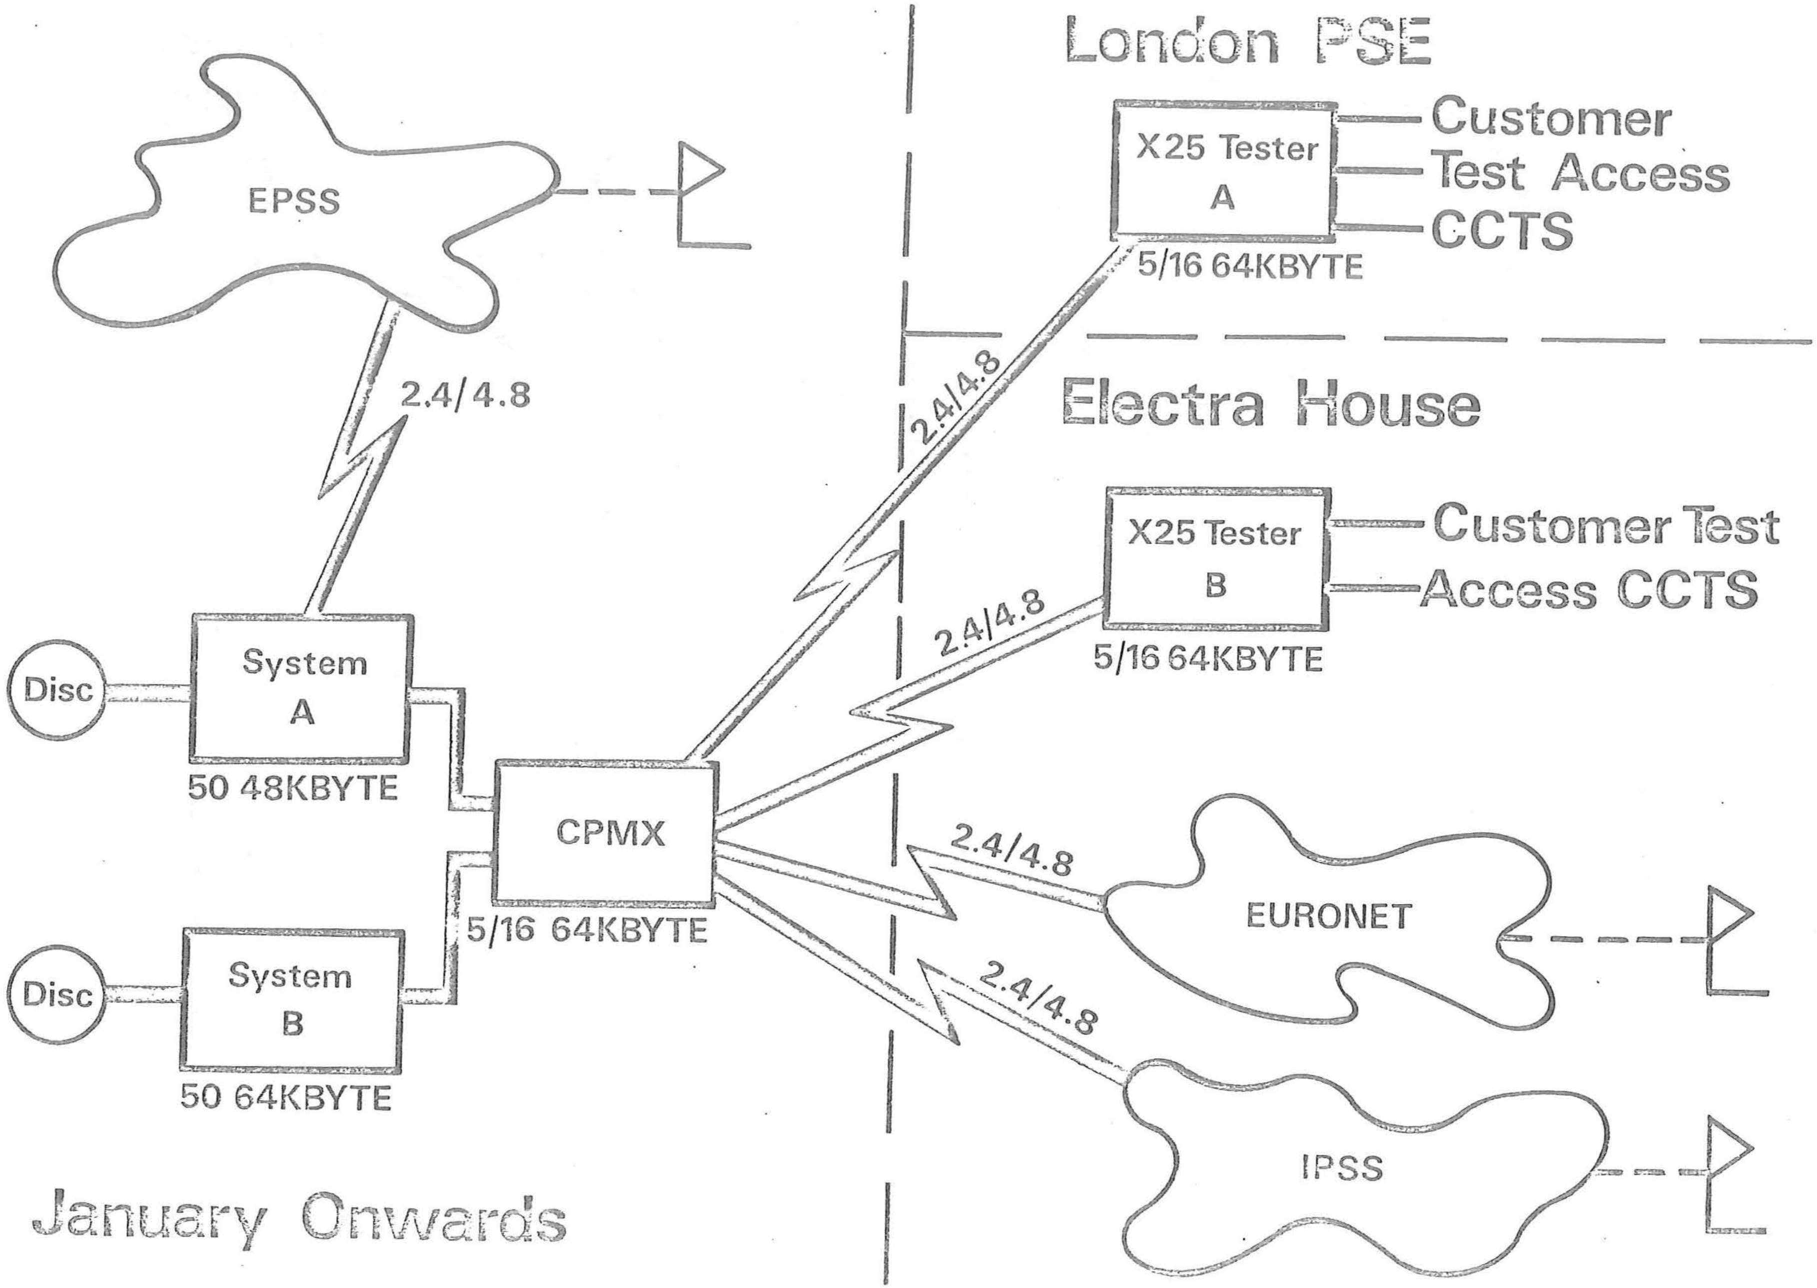

LEVEL 3 COMMANDS

PACKET TYPE (DTE DCE)	COMMAND
Call Request	PX CALL LCN, CD. A, CG. A, FAC. PAIR, . . .
Call Accepted	PX CAA LCN
Clear Request	PX CLR LCN, CAUSE, CODE
Clear CONF	PX CLC LCN
Data	PX D LCN, PR, PS, Q, M, DATA
Interrupt	PX INT LCN, USER DATA
Interrupt CONF	PX INC LCN
RR	PX RR LCN, PR
RNR	PX RNR LCN, PR
REJ	PX REJ LCN, PR
Reset Request	PX RST LCN, CAUSE, CODE
Reset CONF	PX RSC LCN
Restart	PX RES CAUSE, CODE
Restart CONF	PX REC

TRACE FORMATS

LEVEL 2

<L2 TYPE   P/F   N(R)   N(S)   Send frame

>L2 TYPE   P/F   N(R)   N(S)   Receive frame

L2 TYPE:

I
SARM
DISC
UA
CMDR
REJ
RNR
RR

P/F:

P = P bit set
F = F bit set

N(R):

RR
RNR frames
REJ only
I

N(S):

I frames only

TRACE FORMATS

LEVEL 3

L3 TYPE GLCN P(R) P(S) XXXX

L3 TYPE:

CLR
CLC
CALL
CAA
RST
RSC
RES
REC
INT
INC
RR
RNR
REJ
D
??– L3 TYPE unknown

GLCN: First 2 bytes of packet

P(R): RR, RNR, REJ, D Packets only

P(S): D packets only

XXXX: Second 2 bytes of packet

TRACE EXAMPLE (1)

DCE	N(R)	N(S)			
S <I	0	0	RES	1000	FB02
R >RR	1				
R >I	1	0	REC	1000	FF00
S >RR	1				

TRACE EXAMPLE (2)

PX CALL 9,234120200022,234120200022,0,0

DCE	N(R)	N(S)	TYPE	LCN	
R >I	6	2	CALL	1009	OBCC
S <RR	3				
S <I	3	6	CALL	1001	OBCC
R >RR	7				
R >I	7	3	CAA	1001	OFOO
S <RR	4				
S <I	4	7	CAA	1009	OFOO
R >RR	0				

TRACE EXAMPLE (3)

DCE	N(R)	N(S)	TYPE	LCN	P(R)	P(S
R >I	2	4	D	1001	4	6
S <I	5	2	RR	1001	7	
R < RR	3					
R >I	5	3	D	1001	4	7
S <I	4	5	D	1001	0	4
R >I	6	4	RR	1001	5	

Network Services

Dr. J. D. Rice
University of Liverpool

NETWORK SERVICES

Dr. John D. Rice

Computer Laboratory, University of Liverpool

ABSTRACT

This paper attempts to provide an understanding of the background to networking in the U.K. Universities. A survey of the computing resources used in both teaching and research is made. Emphasis is placed on the requirements for data capture, storage, processing, transfer and presentation. A model of the computing provision at a typical campus is considered. A rough model for the national scene is then discussed briefly. In conclusion a number of relevant questions are posed to be answered before networking is undertaken.

INTRODUCTION

The primary aim of this paper is the derivation of a model of the computing resources available to the U.K. universities and the access provided to these resources. One of the principal components of the model is a sub-model of the computing scene on a typical university campus. This is derived from an examination of the types of resource present on and accessed from such campuses. Of particular interest are the movements of data between resource sources and sinks.

There is a great temptation in discussing this subject to succumb to a mere cataloguing of types of equipment and the hardware connections deemed desirable between them. The present author fell into this trap immediately by commencing the first draft of this paper with the sentence "The primary aim of this paper is the derivation of a model of the computing equipment available to the U.K. universities and the inter-connection of such equipment"! The term "computing resource" used in this paper is part of a conscious attempt to escape from such hardware orientation. By a computing resource is meant a commodity, hardware or software, localised or distributed, consumable or permanent, which may be provided as part of a computer service. Examples of such resources are programs, peripherals, data bases and processors. Both resources and their inter-connection form the subject of much of what follows. To complete the definition of terms, it will be sufficient for present purposes to define the function of a "communication network" as the inter-connection of computing equipment and the function of a "computer network" as the inter-connection of computing resources. The former is treated elsewhere at Networkshop 3 in a paper on 'Campus Transmission and Switching'. An associated subject for study is the connection of equipment to a communication network and the software and hardware implications of one aspect of this, 'Connecting Mainframes to a Network', are also to be covered elsewhere. On the other hand, given the requirement for computer networks, it is necessary to formulate abstract models of resource types and formalise communication between the processes controlling and accessing such resources by means of protocol definitions. In later papers the formalisation of a basic 'Transport Service' for such protocols will be discussed, and a selection of protocols for Terminal Handling, File Transfer and Job Transfer will be studied.

As a preliminary to such studies the present paper describes the environment from which the requirement for network services in U.K. universities is derived.

The requirement for computing resources stems from both of the main activities of a university: teaching and research.

TEACHING

Teaching with a need for computer resources can be subdivided into three main areas:

a) subject teaching assisted by computer

b) teaching (and use) of applicable computing

c) teaching of computer science.

These requirements are currently satisfied in a number of ways.

Subject teaching by computer assisted instruction (CAI) or computer managed instruction (CMI) has tended, though not exclusively, to be undertaken on a specialist system ("resource package"), either provided as a central service, or on a smaller scale within a department. When such provision is made, the resource package used is likely to include a variety of special data capture (e.g. special keyboards) and presentation devices (e.g. graphics displays) together with large capacity, fast access, high transfer rate data storage devices. More traditionally such teaching has been in a batch environment on a central computing facility to which access provided from the teaching department is highly desirable. Recently the advent of central interactive services, with terminal access distributed to subject departments or special "teaching terminal areas", has meant an increase in the use of such systems for CAI purposes.

By "applicable computing" is meant the use of computing resources to process data obtained from experimental work in the teaching of a non-computer science subject. Traditionally this has meant the use of "cafeteria" (fast turnround) FORTRAN facilities to process student-written programs to analyse experimental data. Such facilities have normally been provided either solely at the computer centre or at several data centres around the campus with punched card input and line printer output facilities. The more imaginative use by teachers of packages on central mainframes has increased, ranging from the use of teacher-written programs in a "cafeteria library" for the processing of particular experimental data to the use of large engineering packages such as GENESYS for complete design projects resulting in the graphical display of results.

Many universities have installed central interactive systems dedicated to teaching purposes, either as an alternative to or to complement a cafeteria-style service. The predominant language used on such a system is BASIC and access to such a facility is required both for whole student classes and to individuals in departments, particularly in subject teaching laboratories. The existence of cheap minicomputer systems has led departments in some universities to instal their own teaching systems. More recently the availability of microcomputer systems for little more than the cost of a terminal has encouraged the provision of such systems in teaching laboratories, providing a limited range of package programs.

The teaching of computer science produces a wide spectrum of requirements for computing resources. Since many of these are concerned with program and data structure they are typically satisfied by a departmental system, whose chief selection criteria may concern the hardware and software architecture of the system and the language processors it supports. Central computer service machines rarely seem to meet such criteria! If comparative studies of computer systems are to be undertaken, then there will be a requirement to access facilities outside the local campus. Low volume of traffic and frequency of access may be handled via dial-up interactive or remote job entry (RJE) facilities; higher volume of traffic and frequency of access is only likely to be provided for teaching purposes to remote resources to which permanent connections already exist.

To summarise, the types of resources which are involved in the satisfaction of teaching requirements are:

a) CAI systems
   - special input devices
   - graphical presentation devices
   - high performance/capacity mass storage devices

b) Batch processing systems

   - cafeteria services
   - specialist packages
   - RJE facilities

c) Interactive services

   - general purpose
   - departmental special purpose

d) Experimental computer science systems

   - structured architecture
   - language processors

e) Remote facilities

   - off the campus.

The above list is not exhaustive by conveys a picture of the range of resources required.

RESEARCH

It is impossible within the scope of the present paper to comprehensively cover the range of research activities which generate a requirement for computing resources. Indeed, a set of requirements will be stated in general terms and relevant examples of each will be quoted. The following areas of interest may be distinguished:

a) data capture
b) data storage
c) data processing
d) data transfer
e) information presentation.

In any particular research programme each area will be present, though one will probably be dominant. Therefore in the following discussion each area will be considered in turn but particular examples will be completely dealt with under the heading of the dominant area.

Data Capture

In social and economic science departments much research involves the collection and analysis of survey data. Traditionally such data has been transferred manually from survey form to punched cards, but increasing use is now made of data capture

direct from machine-readable survey forms using optical character recognition (OCR) or optical mark recognition (OMR) equipment. The cost of such equipment and its pattern of use tends to favour its location centrally.

In science, engineering and medical departments much data is captured directly from experimental rigs. The user may require:

a) data logging for remote processing
b) local data reduction with subsequent data logging for remote processing or
c) data logging, processing (local and/or remote) and feedback of results.

The first two requirements imply a need for data transfer between the data capture and data processing sites. The third may also do so, in which case the speed of data transmission in each direction may be significant, along with the response provided by the remote processor.

The power of systems used for local processing varies considerably, from slow micro-processor-based systems to large mainframes. Similarly the peripheral endowment of such "departmental computers" varies from analog and/or digital inputs with paper tape punch output only to large numbers of magnetic tape and/or disk drives.

The media used to log data for remote processing range from paper tape, through cassettes, cartridges and floppy disks, to magnetic tapes and disks. Some data is logged by the emulation of an RJE terminal. Some equipment produces charts which require digitising to convert data to machine-readable form.

For many research projects however, data is still captured in more traditional manner by keyboard input, either onto a machine-readable medium such as paper tape or punched cards, or directly onto a magnetic medium. Online direct data entry may either be provided on a dedicated system or as one of the services on a data processing system.

Data Storage

Great advances are taking place in storage technology. but in reviewing the current general situation in U.K. universities it is sufficient to assume that the majority of data is stored on magnetic media either instantly available for processing (on-line) or archived (off-line). Data is normally held online at one of three locations:

a) in association with a central data processing facility e.g. main university computer service mainframe

b) independently e.g. campus network filestore

c) in association with a localised facility e.g. a departmental computer system.

In the above it is not necessary to distinguish between simple file systems and data base systems. However, it should be noted that increasing use is being made in research of specialised data bases, the majority of which are maintanined by commercial organisations at sites remote from university campuses, in many cases overseas, particularly in the U.S.A. and Europe.

Data Processing

Much data processing in universities is concerned with user-written programs and user-captured data. This processing will normally take place on the most suitable computer system available to the researcher. The only proper definition of 'suitable' in this context must be 'that which most positively aids the progress of research'. Unfortunately in many (most?) cases 'suitable' is usually interpreted as 'cheapest' or 'most accessible'.

Two particular types of data processing demand special attention. Many projects involve the manipulation of exisiting data. In such cases the major requirement is a facility with hardware and software adapted to the efficient processing of the type of data involved, which may mean a 'data base' environment.

In other research projects the emphasis is on the use of special-purpose package programs to analyse experimentally-obtained data. In such cases the availability of an appropriate environment for the package is the over-riding consideration. Such an environment may be related, for example, to the number of significant figures provided for numerical computation or to the provision of high-speed vector processing facilities.

It is obvious that not all the facilities required by the researchers in any particular university will be located on that university's campus; but the need for access to all facilities from each campus is equally clear.

Data Transfer

This is primarily concerned with the requirement to associate the location of the data processing resource with that of the data storage resource. The requirement may be conveniently discussed under the two headings:

a) bulk data transfer
b) transaction-based data transfer.

Bulk data transfer

The simplest example of the bulk data transfer requirement is that of the research worker with data stored on punched cards which he wishes to have processed by a specialised package. If the package is available on the local computer centre's mainframe, then the data will normally be transferred from an RJE terminal near the user's department to his local machine for processing, and any associated results returned to the same location. If the package is located, say, only at a national centre such as UMRCC, then the data requires transfer over a longer distance, probably from a single RJE terminal dedicated to this service.

If the data is captured on a local storage system of some type (e.g. departmental computer, key-to-disk system) then a similar need exists as above, but the satisfaction of this need may present more complications. If data is only destined for a local mainframe, then the connection of the equipment in RJE emulation mode is a common solution. If the data is destined only for a remote mainframe, then a similar type of connection may be used. Typically such connections may be to a

specialist, e.g. SRC, remote facility. If data made be destined for different mainframes from one local source, then at the least an RJE-switching facility is required.

On a larger scale, the local storage system may be the local computer centre's mainframe system and the transfer of partially-processed data from such a facility to be processed at a remote resource presents larger problems. Such a requirement may be met by RJE emulation, as in the use of the SWAN system in ICL 7903 communication processors to connect ICL 1900 processors, e.g. Birmingham's 1906A with UMRCC's 1906A.

Transaction-based data transfer

Normally there are not stringent requirements on the rate at which bulk data transfers should occur, although these may be dictated by the volume of data involved.

Some research projects, however, require fast selective access to remotely-held data. The lack of provision for such access in the past has tended to lead to the storage locally of selected parts of remote data bases.

Information Presentation

The advent of the microprocessor, and particularly its introduction into peripheral equipment, has begun to revolutionise the whole way in which the presentation of the results of data processing is handled. Increased intelligence in the local presentation resource is beginning to change, for example, the method by which a graphical representation of data is provided by a mainframe to a graphical terminal. In the past there has been a straightforward bulk data transfer of elementary vector plotting instructions but it is now possible with new graphics terminals to conduct a transaction-based dialogue with commands at a picture-element manipulation level.

The majority of results output in U.K. universities are produced in textual form on line printers. On large campuses several bulk printing terminals are provided , usually in association with a data input device as part of an RJE terminal. Where large quantities of textual information are required for long-term storage and retrieval many sites are now offering facilities for output on microfilm or microfiche. In the past the majority of interactive terminals have been slow Teletype-like devices. The trend in recent years has been to replace these by VDUs (Teletype-compatible or with local editing features) often associated with a hard-copy device (e.g. ICL 7502 systems, Ferranti PT7).

In the future it is likely that medium speed hard-copy printer/plotter keyboard terminals will become increasingly common place, replacing distributed line printers, and that graphics/alphanumeric VDUs with local data editing and processing capabilities will replace existing interactive terminals.

At present a number of universities have small packaged systems dedicated to the manipulation and display of graphical data. Such systems will clearly remain, supporting increasingly sophisticated presentation devices; though it is likely that the advent of cheap reasonable-quality graphics presentation devices will mean that the requirement for distributed access to such systems will grow.

There is also a requirement for the presentation of information in machine-readable form (e.g. for numerically-controlled machine tools), and there is an increasing need to produce output on cheap magnetic media (floppy disk, data cartridge) for such purposes. In this connection some universities provide media conversion services.

A MODEL CAMPUS

From the foregoing discussion we can derive a model of a typical university campus's computing resources:

1. Central services

   (a) mainframe (batch) system
   - general purpose computing facilities
   - "cafeteria" service
   - specialised packages

   (b) interactive system
   - conversational languages
   - specialised packages

   (c) CAI system

   (d) graphics system
   - special peripherals
   - specialised software

   (e) media conversion service
   - magnetic media
   - OCR,OMR
   - digitiser

   (f) online data entry service.

2. Departmental services

   (a) data logging system
   (b) process control system
   (c) CAI system
   (d) experimental computer science system
   (e) general data processing system.

3. Access services

   (a) data base services

   (b) specific problem-oriented services
   - machine-architecture
   - calculation accuracy
   - data manipulation facilities
   - mill power
   - main storage capacity
   - subject-dependence
   - e.g. SRC nuclear physics services

   (c) other non-local services
   - e.g. machine dependent packages.

Access services may be provided by:

(a) remote data entry terminals
(b) interactive terminals

clustered centrally on the campus, clustered in several locations on the campus, or located in individual departments. It should be noted at this point that lack of financial resources has meant in the past that some universities have provided few if any local services and have become greatly dependent on non-local resources to meet local computing requirements. Such remote services may be provided:

a) by a national centre (e.g. UMRCC, NUMAC)
b) by another regional site (e.g. database services in the North-West at Liverpool)
c) by a Research Council (e.g. SRC, NERC)
d) by commercial bureaux in the U.K.
e) by overseas installations via international networks.

Inter-Connection

The aim of each university must be to provide each of its members with access to the computer resources required for the optimal satisfaction of the overall requirements of teaching and research. Many constraints, financial and technical, ensure that this goal is not attained by all.

If the model campus is to bear any relation to reality, then it must be recognised that not all that is technically feasible will be economically viable, nor in some cases can technical solutions to requirements be provided, even though finance may be available.

As a general principle, access is required from most parts of a campus to central services. A distribution of terminal equipment is therefore required according to needs, which may include the requirement for clusters of terminals, both for teaching classes and for research groups. One of the functions of the central computer service is to anticipate future needs, so distribution technology needs to take account of this factor.

Economies may be effected by providing access to several services from a single terminal. The question then arises as to whether a single user accesses different services from this terminal or several users each access a single service. In the latter case, no inter-connection between services may be required. In the former case, data transfer between central services is likely to become a major requirement.

To cover both eventualities the model campus may support both terminal switching and central resource intercommunication.

The review of resources has already indicated that terminal functions are currently being performed by equipment of varying levels of sophistication, ranging from simple VDUs to large mini-computer systems. Since the trend is towards more complex terminals, with the inclusion of microprocessors, the campus terminal switch needs to support the whole range of terminal complexities, but in a standardised way.

Experience indicates that full use of remote resources is only obtained when access is provided at a level comparable to that provided to local resources. Duplication of access distribution points (terminals) is clearly not viable, so the model campus network must be capable of full connection to remote services.

NATIONAL MODEL

The Computer Board has created a concept of Regional Groups of universities and funding policy has, to a greater or lesser, welcomed or unwelcomed, extent created a certain amount of intra-regional dependence in the provision of computer services. It is natural therefore to see a region as a collection of model campuses, each with particular specialisations (e.g. interactive packages, data bases etc.). Outside of the regions lie the national centres (Computer Board and Research Councils funded) and "the world". It is reasonable to assume that the model campus provides access to a limited range of extra-regional services. A regional network therefore provides, almost as a by-product, a rationalisation of such access on a regional basis. The physical realisation of such a conceptual model will depend on traffic volumes and desired responses for cross-network interactions.

CONCLUSIONS

No two universities are alike (comparison with the model campus will clearly demonstrate this), however there is a sufficient commonality in pusuits that a general model is not without interest as a basis for the analysis of an individual site's requirements.

In approaching the subject of network services on a departmental, campus, regional, national or international level, the following questions may be posed:

1. What resources are provided and required in the area under consideration?

2. Why is a connection required between A and B?

3. What are the constraints imposed on establishing such a connection?

4. What use will be made of the connection?

5. What side effects are there of its establishment?

Only when satisfactory answers have been obtained can investigation properly commence into both the hardware and software aspects of networking.

Report on Attendance at Computer Communications Conference (5th-8th Sept) and Visit to the National Bureau of Standards (11th Sept) in Washington D.C., USA

R. Chisholm
Edinburgh Regional Computing Centre

NETWORKSHOP 3

REPORT ON ATTENDANCE AT COMPUTER COMMUNICATIONS CONFERENCE (5th-8th SEPT) AND VISIT TO THE NATIONAL BUREAU OF STANDARDS, (11th SEPT) IN WASHINGTON D.C. USA

Since I was actively engaging in an evaluation of current trends and practices in local area networking the IEEE sponsored conference on "Computer Communications Networks - Compcon 78" provided me with an opportunity to hear and meet some of the people active in this field.

The conference was held in the Capitol Hilton Hotel, Washington D.C. on the 5th-8th Sept. 1978. The proceedings were organised in two parts - the 5th Sept. being devoted to three parallel pre-conference tutorials - the 6th-8th Sept. was devoted to the presentation of selected papers. Appendix A contains the conference program.

Mon.4th Sept.

On registering I was presented with a book containing a complete set of the conference papers and a book covering the tutorial session I had chosen to attend.

Tues.5th Sept.

Of the three parallel tutorials I decided to opt for "A Practical View of Computer Communications Protocols" presented by J.M. McQuillan - Bolt, Beranek & Newman Inc. McQuillan is one of the designers of the ARPA network and not unnaturally his presentation leant heavily on his experiences with this network.
The morning session was given over to an evaluation of the various influencing factors in network design ranging from the types of transmission facilities through to the choice of subnet protocols.
The afternoon session dealt with a design study and performance analysis of a host-host protocol and terminated with a review of some of the currently implemented networks, both national and local. Although I was fairly familiar with some of the topic areas I was very impressed by the quality of the content and the overall presentation.

Wed.6th Sept.

The morning session was devoted to the conference opening speeches which were followed by three keynote speeches delivered by representatives of the government, a/

a telephone company , and the computing industry. Whilst much of what was said impacted more on the domestic scene some of the points made will have global significance.

I noted that -

* AT & T, who until 1979 are restricted to offering only common carrier facilities, have an application with the FCC for the provision of a nationwide computing service, in addition to existing offerings. The system they are proposing is known as ACS (Advanced Computing System).

* With the increases in data traffic rates (bits)there is a movement to the use of communications satellites as an economic alternative to land lines. The government are encouraging the development of low cost terrestial links to satisfy this market.

* A lot of use is being made of portable radio equipment as the basis for the development of cheap local area packet broadcast networks.

In the afternoon I attended -

(1) Session 3 - "Flow Control Analysis". Although moderately interesting, none of the three presented papers bore any direct relevance to my area of interest.

(2) Session 7 - "Packet Bus Structures". The first paper described work being done on the development of a high bandwidth parallel communications network designed to interconnect closely associated processors and peripherals.

The second paper detailed an Ethernet type network currently being implemented at the National Bureau of Standards. Resulting from discussion I had with the authors of these papers arrangements were made for me to visit the NBS on the 11th Sept. I report elsewhere in this paper more detail of this network.

The last paper in this session described work being done at Digital Equipment Corp. on a method of slotted contention bus arbitration.

Thurs./

Thurs.7th Sept.

I attended the following sessions -

Session 8 - "Technology"

The first paper described work being done at Sperry Univac on the design of network circuit switching elements; mostly for use in TDM applications. The emphasis being on formalising design procedures.

The second paper dealt with techniques developed to work out optional routings in hybrid packet and circuit networks.

The third paper described a new technique being developed to provide high bandwidth low cost "in house" data communications. This system developed by IBMZürich, makes use of the fact that infra red radiation directed at a wall or ceiling will be diffusely scattered by that surface. Using this technique it is possible for a host computer to communicate with a cluster of terminals located in the same room.

Session 12 - "Micronetworks"

The first paper described a local area network which had been modelled on the ARPA network. The IMP's are based on proprierty microcomputers. From the presentation it appeared that although the design goals were modest, a very useful and cost effective network had been developed.

The second paper described work being carried out on the development of a high bandwidth ring network. The technique used for transmission was "time-slotting". Because of the high data rates used (1.544 MHZ) this system contained a lot of special purpose hardware.

The third paper described work being done on the design of multi-microprocessor networks using Multi-Tree-Structured graphs.

Session 13 - "Network Architectives"

The first paper presented was on the BANCS network developed at Bell Laboratories. This network currently supports traffic from 800 terminals all accessing a large IBM mainframe. The projections are that by 1980 3,400/

3,400 terminals will have access through this network to six large host mainframes.

The second paper described a high bandwidth inter-computer communications facility developed by RCA. The heart of the system is a "switch" capable of interconnecting up to 255 computers, which can support a load of 74M bits/sec. The links between the switch and the processors run at 3M bits/sec using the ADCCP protocol.

The third paper, although orientated towards large process control systems, described design work that had been done on an Ethernet type network which was constructed as a rectangular matrix, with each node having a connection to two separate buses.

Session 19 - "Network Operating Systems"
The first paper was given by the research director of Prime Computers. He described a system which was being constructed by interconnecting computer cells. These cells are interconnected by a high speed packet type ring network to form a multiprocessor system.
The following three papers described work being undertaken at various organisations to augment and improve the data communications facilities available under the UNIX operating system.

Session 24 - "Hardware"
The first paper, from the University of Delaware, described a facility they have designed to assist in the development of hardware and software support for micro-computer systems.

The second paper described a logic circuit which has been developed to overcome the effects of random "meta-stable" states which flip-flops occasionally display when improperly triggered.

The last paper covered various hardware/software problems the author has encountered when implementing systems using LSI circuits. The basic message/

message was "don't assume that data sheets are either accurate or exhaustive".
The problem lies in the increasing complexity of devices which are becoming more difficult to accurately parameterise and test. He gave some examples to illustrate the point - we could add some to his list!

Session 27 - "The Network Design Process- An Integrated Approach"
The four papers presented were all from the same stable - the Network Analysis Corp. They dealt in some detail with the techniques which they have developed and use to turn customer requirements into actual network designs.

<u>Mon.11th Sept.</u>
As a result of discussions with the authors of a paper I was invited to visit the Institute of Computer Science and Technology at the National Bureau of Standards, to see and discuss the work they are doing in local area network.
They have developed a 1M bit/sec co-axial cable ethernet type packet broadcast system for use within their department. They have opted for a hardware oriented solution - all the bus access management and error recovery is handled by dedicated logic thus providing the user with a fairly simple and manageable interface; which can be bit serial or byte parallel. They chose this design philosophy because they felt it would allow them to attach a much broader spectrum of devices.
During our discussion I indicated that we might be interested in acquiring, in some form or other, the system technology they have developed; they appeared fairly enthusiastic.

<u>SUMMARY</u>
On my return to this country I wrote to the NBS confirming my interest in their system and asked them to indicate how best we proceed, should we decide that their system would satisfy our requirements.
One thing that is obvious from my talking to people in the States - there is no standardisation in the area of local networks.

R. Cushell
29/9/78.

Connecting Mainframes to Networks

Dr. P. S. Kummer
Daresbury Laboratory

DARESBURY LABORATORY

COMPUTER NETWORKS DEVELOPMENT GROUP

CONNECTING MAINFRAMES TO NETWORKS

P S Kummer

25 September 1978

This paper considers ways in which mainframe computers can be connected to a communications network. Firstly, a definition of a mainframe is given which, although differing from that normally used, is useful for the following discussion. The type of network considered is also restricted to a public, packet switched service.

Both the hardware and software aspects of the connection are considered and examples are taken from implementations at Daresbury Laboratory. Although these implementations are based on EPSS-like protocols, the use of X25 would not significantly change the discussion.

The following discussions are not meant to be exhaustive but are intended to highlight some of the more interesting aspects out of a real implementation.

What is a Mainframe?

The advent of more and more powerful "mini" and "midi" computers means that these machines are now able to offer services that have previously been limited to the larger machines. Thus the term "mainframe", implying large computer, is restrictive when discussing connection to a computer network. In this paper I will expand the definition to include all machines providing a service to the users of the network. This includes:

Small machines providing simple services (e.g. a PDP11/05-based network status monitor capable of providing status information, on request, to an interactive terminal.)

Medium machines providing a reasonable amount of processing power (e.g. a local editing and batch processing service on a GEC 4080.)

Large machines providing powerful services (e.g. an IBM 370/165 providing batch and interactive services with a large on-line file store.)

Very large machines providing specialised services (e.g. CRAY, DAP).

The last type of machine usually requires a relatively large machine acting as a front-end to provide job collection and outputting facilities. Thus the very large machines will not be considered in what follows.

What is a Network?

The only type of network considered here is a public network based on non-proprietry standards (i.e. X25). Manufacturers' networks (DECNET, SNA, etc.) will not be discussed.

Hardware

Small and medium size machines may usually be connected to a public network using manufacture supplied hardware. This is because these manufacturers are more conditioned to connection to other manufacturers hardware. HDLC interfaces are being provided by some manufacturers (e.g. DEC, GEC).

[Daresbury uses the CAMAC standard to interface peripherals to computers. As this standard is computer independent only one HDLC interface needs to be designed.]

One of the limiting factors with manufacturer supplied communications interfaces is the maximum data rate supplied. The best that can be obtained is:

48 Kbits/sec - synchronous

9.6 Kbits/sec - asynchronous.

Note that for computer-computer communications, asynchronous becomes synchronous as far as the receiver is concerned.

The use of synchronous transmission at speeds greater than 9.6 Kbits/sec may require that a DMA interface of some sort is used. (This is because the interrupt response time of most systems is too long to guarantee being able to avoid "overrun" conditions.) This will increase the cost of the interface by a significant amount.

If we consider a CAMPUS network then the use of standard synchronous interfaces may not be ideal for two reasons:

(a) Some form of MODEM or MODEM eliminator is required, thus increasing the cost.

(b) The maximum data transfer rate of 48 Kbits/sec may be too low.

Note that one of the properties of a store-and-forward network is that each additional link in the communications path reduces the effective end-to-end data rate, if multiple buffering cannot be used (as is the case for interactive terminals).

Daresbury uses a fast serial link module, on site, capable of running at 5 Mbits/sec, with a hardware handshake which allows the receiver to control the actual data rate.

Under program control (no DMA) a PDP11/05 can drive this link at about 160 Kbits/sec whereas the corresponding value for a synchronous link is about 10 Kbits/sec.

Large machines usually require some form of front-end processor. This is because the required hardware interfaces are not available (and would probably be expensive if they were!). Manufacturer supplied front-ends are designed for efficient implementation of that manufacturers own communications protocol and are not necessarily very good for interfacing to the public network. Again, it is possible that the required interfaces to the public network are not available for the front-end processor.

An alternative solution is to use a medium size machine (perhaps from another manufacturer) as a front-end processor. The communications hardware requirements for this processor may be split into two parts:

(a) Interface to the public network.
- available on many machines as described above.

(b) Interface to the large machine.
- more difficult, reduces the range of machines available.

Software

For small and medium size machines there is a direct hardware connection to the network and this will be reflected in the software as shown in figure 1. The only comment I would like to make is that it is worthwhile combining the transport service and switching into one machine. Besides reducing overall system cost, this provides extra flexibility as services may be moved between machines with the minimum of effort. This configuration is shown in figure 2.

For large machines the situation is not so clear and it is possible that a universal solution does not exist. The following points help to complicate the situation.

(a) The nature of the manufacturers operating system and the availability of documentation.

(b) The number and nature of local modifications.

(c) The expertise of the system programmers.

(d) The willingness of management to develop and run a relatively complex piece of software.

(e) The required data traffic characteristics.

(f) The money available for new hardware.

(g) The range of large machines around and the differences within a single range.

(h) The number of sites with similar machines/requirements who are willing to cooperate with each other.

As it is impossible to describe all of the possible solutions, I will describe the one I know best which is the one at Daresbury Laboratory.

The Daresbury Laboratory Solution

In what follows it is important to remember that this is an "evolved" system where the present situation was not visable during the initial design stage ($\sim$ 1970).

The first solution is shown in figure 3 and was designed to interface a non-network system to the central IBM 370/165. At the time of this solution considerable expertise existed within the 370 systems group. Also, the only hardware interface available was a single address interface necessitating logical channel multiplexing over the link. The solution adopted presented a standard IBM device software interface to the application programs (including "system" software such as operator console support).

When the IBM 1802 front-end was replaced by an Interdata 85 computer it was decided to introduce a packet switched network for connection to small machines on and off site. At the time very little software development effort was available on the 370 but a considerable amount of effort was dedicated to the Interdata. Thus the support on the 370 was only changed minimally and all knowledge of the network was confined to the Interdata.

This produced the situation shown in figure 4 where the 370 and Interdata software conspired to make devices on the network look like standard IBM devices. In particular, EPSS ITP terminals were mapped into IBM 2260s, thus ensuring that no modifications were required to IBMs TCAM and TSO programs.

The great advantage of this approach is that standard programs written to drive standard IBM devices may be used without change when the device is remote and connected via the network.

The disadvantage of this approach is that some applications in the 370 may be improved if they have knowledge of, and control of, their network communications channels. This was provided at an early stage for some applications by an extra protocol above the multiplexed logical channel but, in order to simplify the applications, the protocol provided minimal virtual call control.

The first application to be restricted by the "remoteness" of the network was the time-sharing service (TSO). The solution here was to provide full network call handling facilities in a new communication process (replacing IBMs TCAM). In this case the Interdata acts only as a switch using a pair of logical channels (to provide full duplex) as a communication link to the TSO transport service. This is the situation as it exists at present and is shown in figure 5.

Experience of this system over the past two years has shown that the interface at the device level is generally no longer required and that a more appropriate interface level is at the device dependent, logical I/O level (i.e. at the logical GET/PUT level rather than the "start channel program" level). Work has recently started on providing a transport service at this level and will produce the system shown in figure 6. This leaves the IBM console support as the only application to not use the transport service. It is hoped that the components of the system will be able to be switched around to give the system shown in figure 7 which we consider to be the ideal way of interfacing the IBM to the network.

(Note: The interface between the IBM and the Interdata will use multiple addresses, allowing the Interdata to emulate one input and one output device to provide the full duplex link required.)

Some Final Remarks

1. It is good policy to make the interface between the application level and the transport level the same as that between applications and the mainframe's device independent I/O system.

2. It is useful for the application to have full control over network virtual calls. Rather than develop another protocol with all the facilities of the standard network protocol (e.g. X25) then the standard network protocol should be continued right into the mainframe.

3. Standard manufacturer supplied communications hardware may prove limiting in data transfer rate, especially for CAMPUS networks. This will be most obvious for interactive terminal response when there are many data links in the communications path. However, the figures shown in figure 8 may be of interest as they show the low average data rate required to support a large mainframe.

FIGURE 1. SMALL/MEDIUM MACHINE CONNECTION

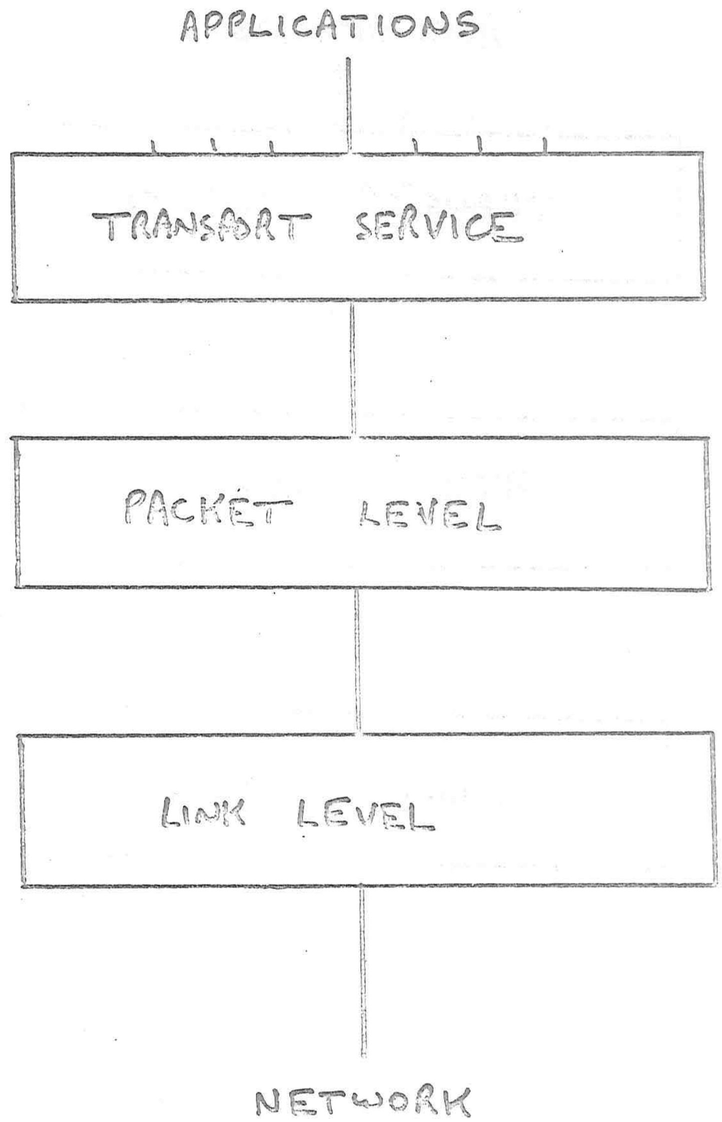

FIGURE 2. TRANSPORT SERVICE AND SWITCHING IN SAME MACHINE.

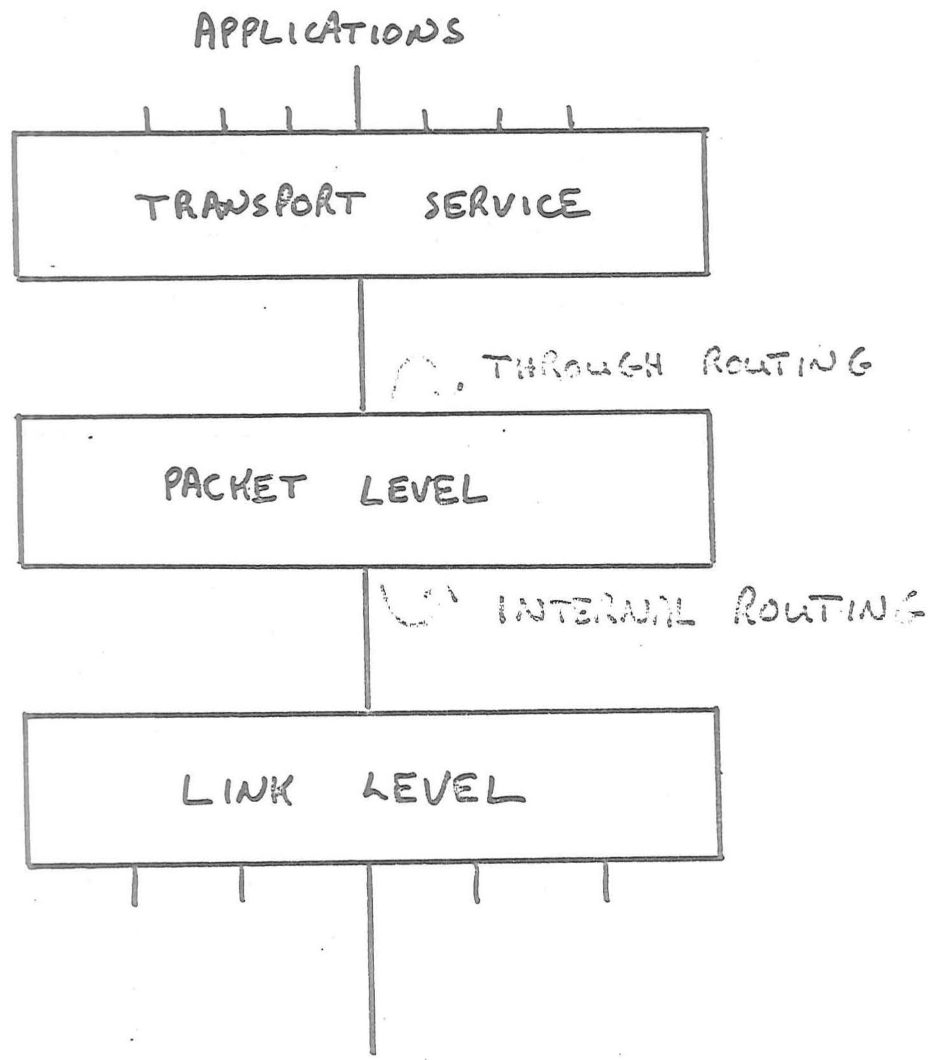

FIGURE 3. D.L. FRONT-END SYSTEM. ~1974

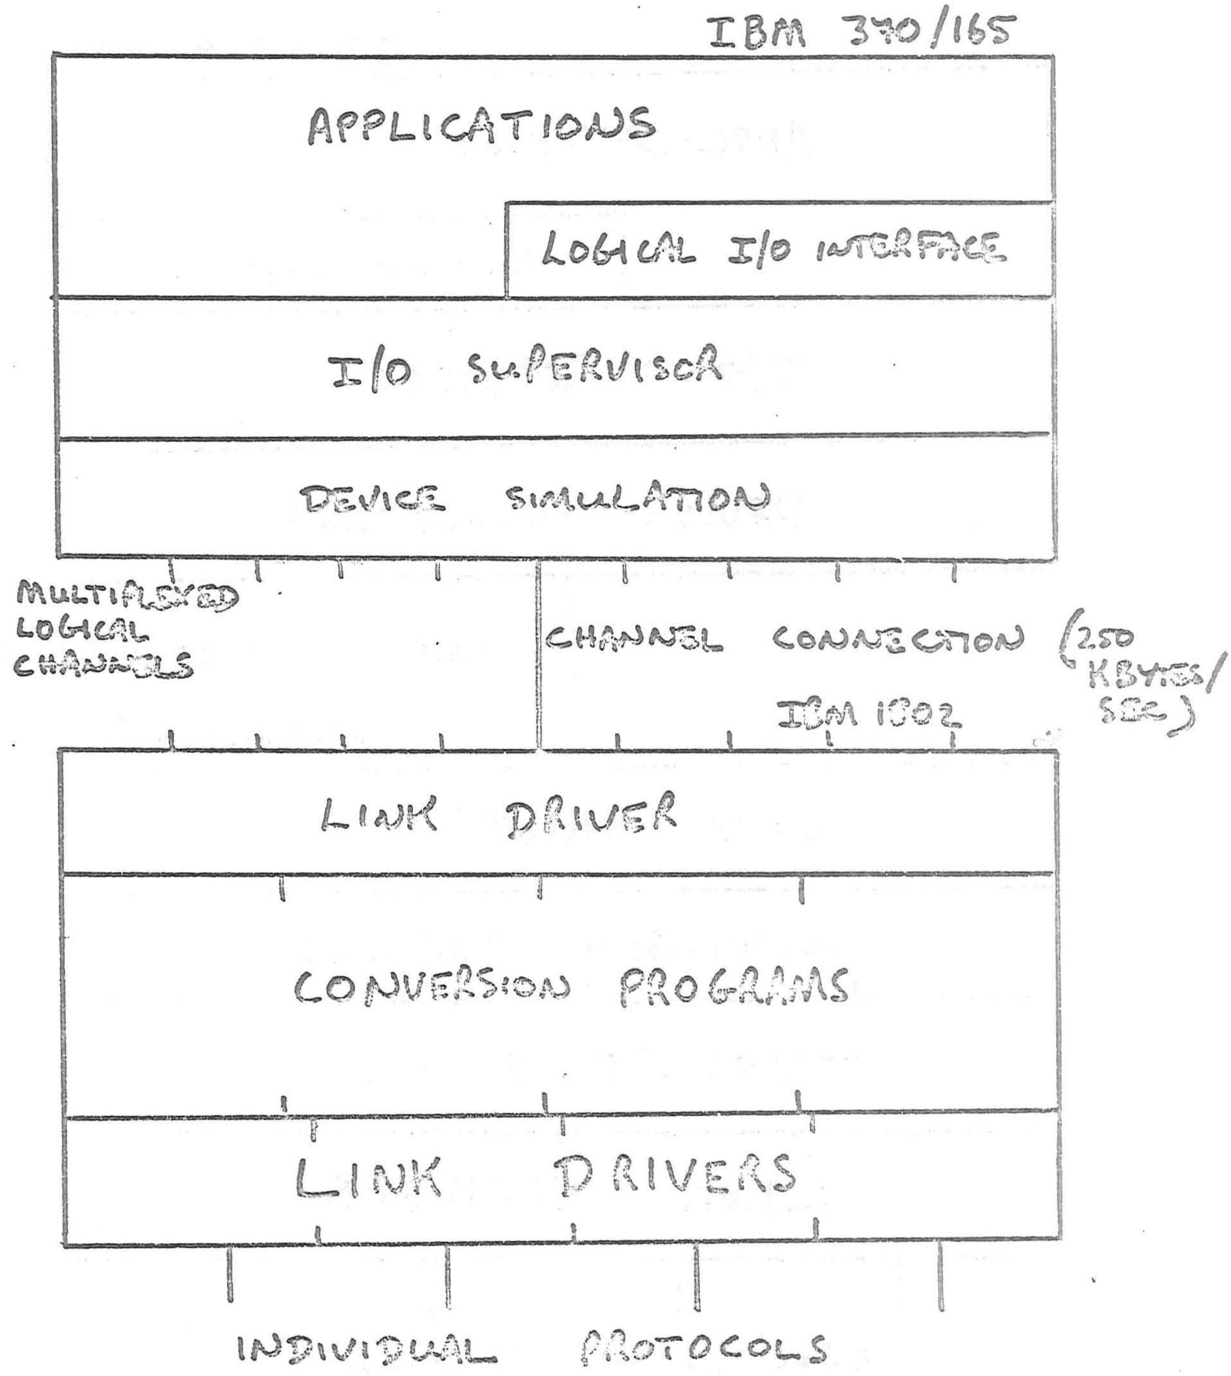

FIGURE 4. D.L. FRONT-END SYSTEM. ~ 1976

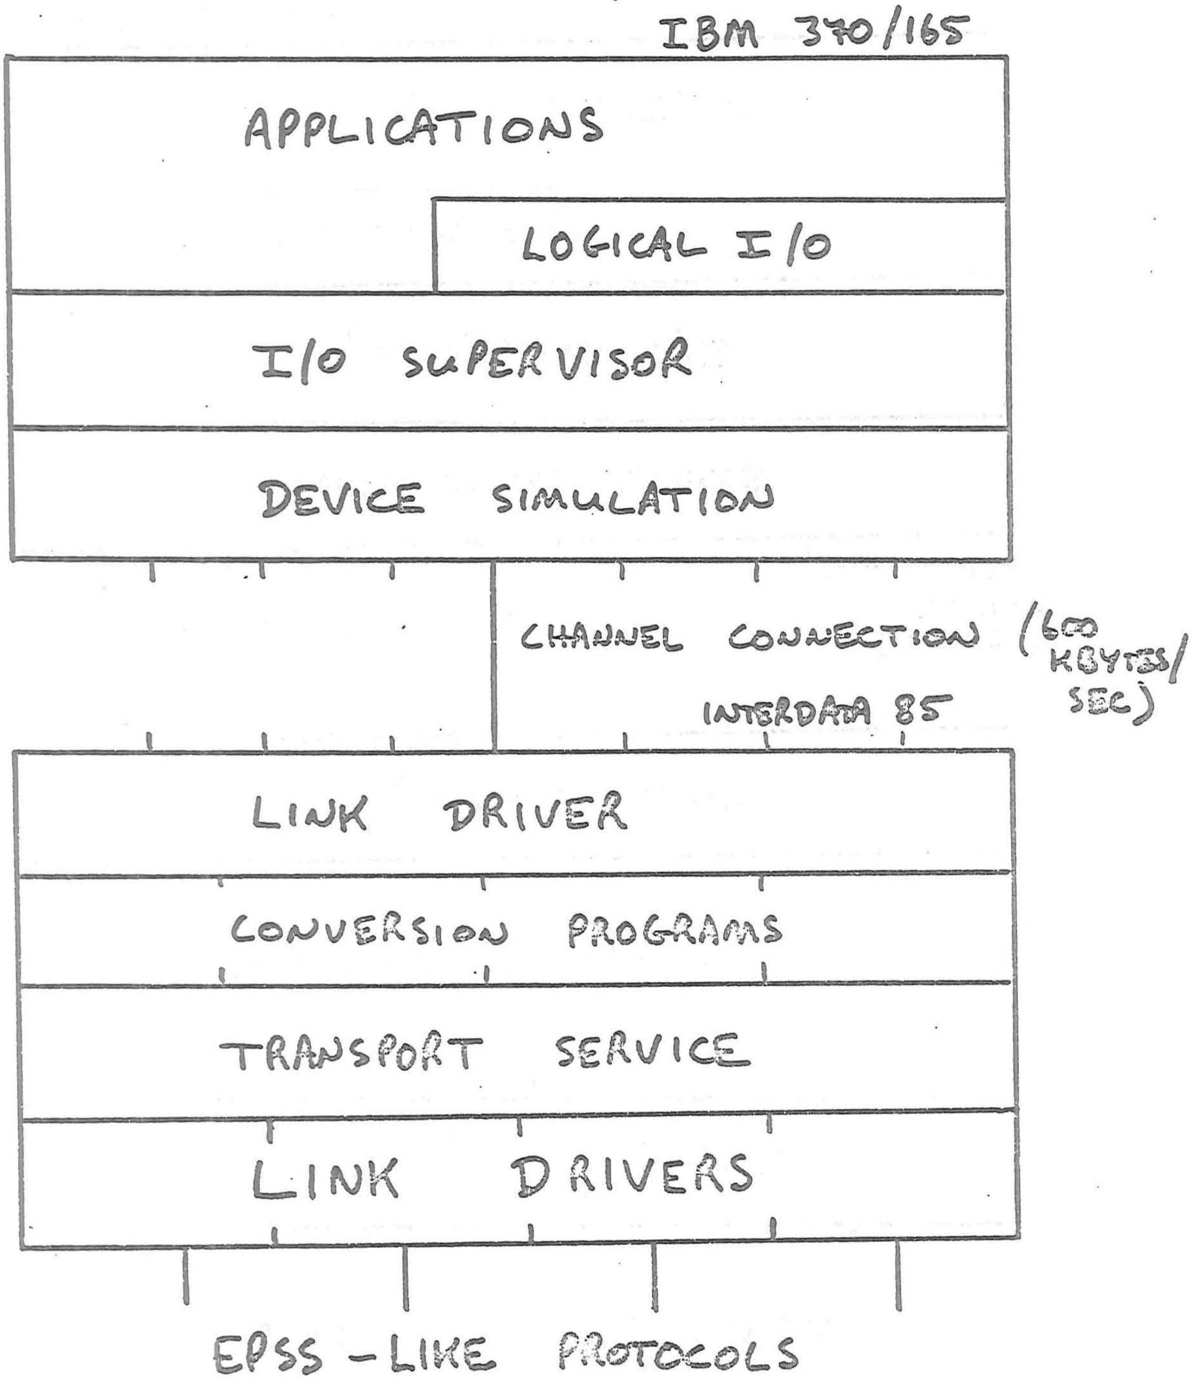

FIGURE 5. D.L. FRONT-END SYSTEM. - 1978

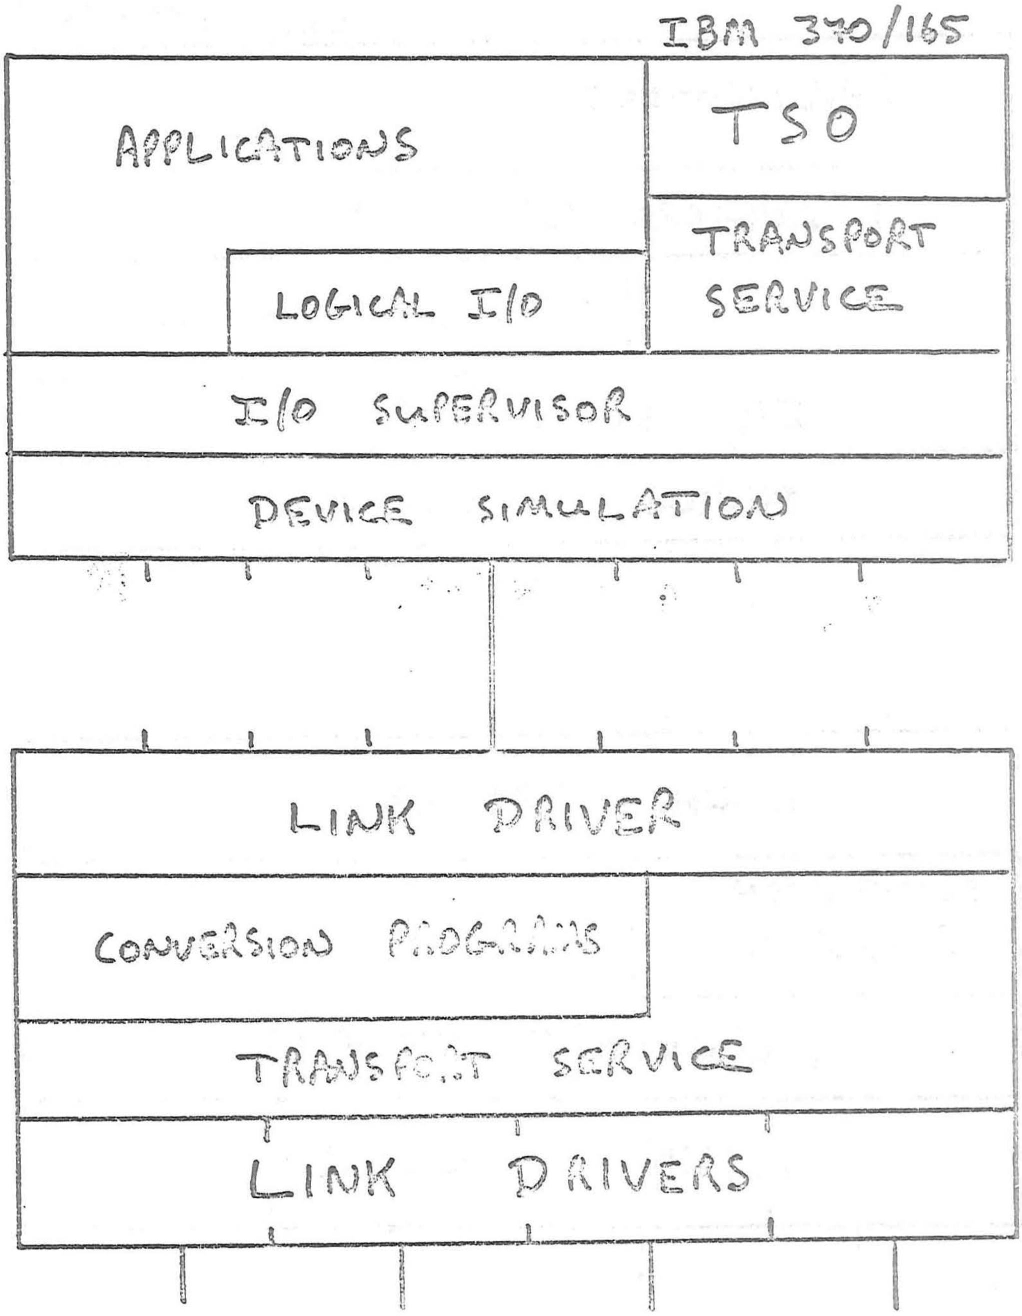

FIGURE 6. D.L. FRONT-END SYSTEM. ~1979

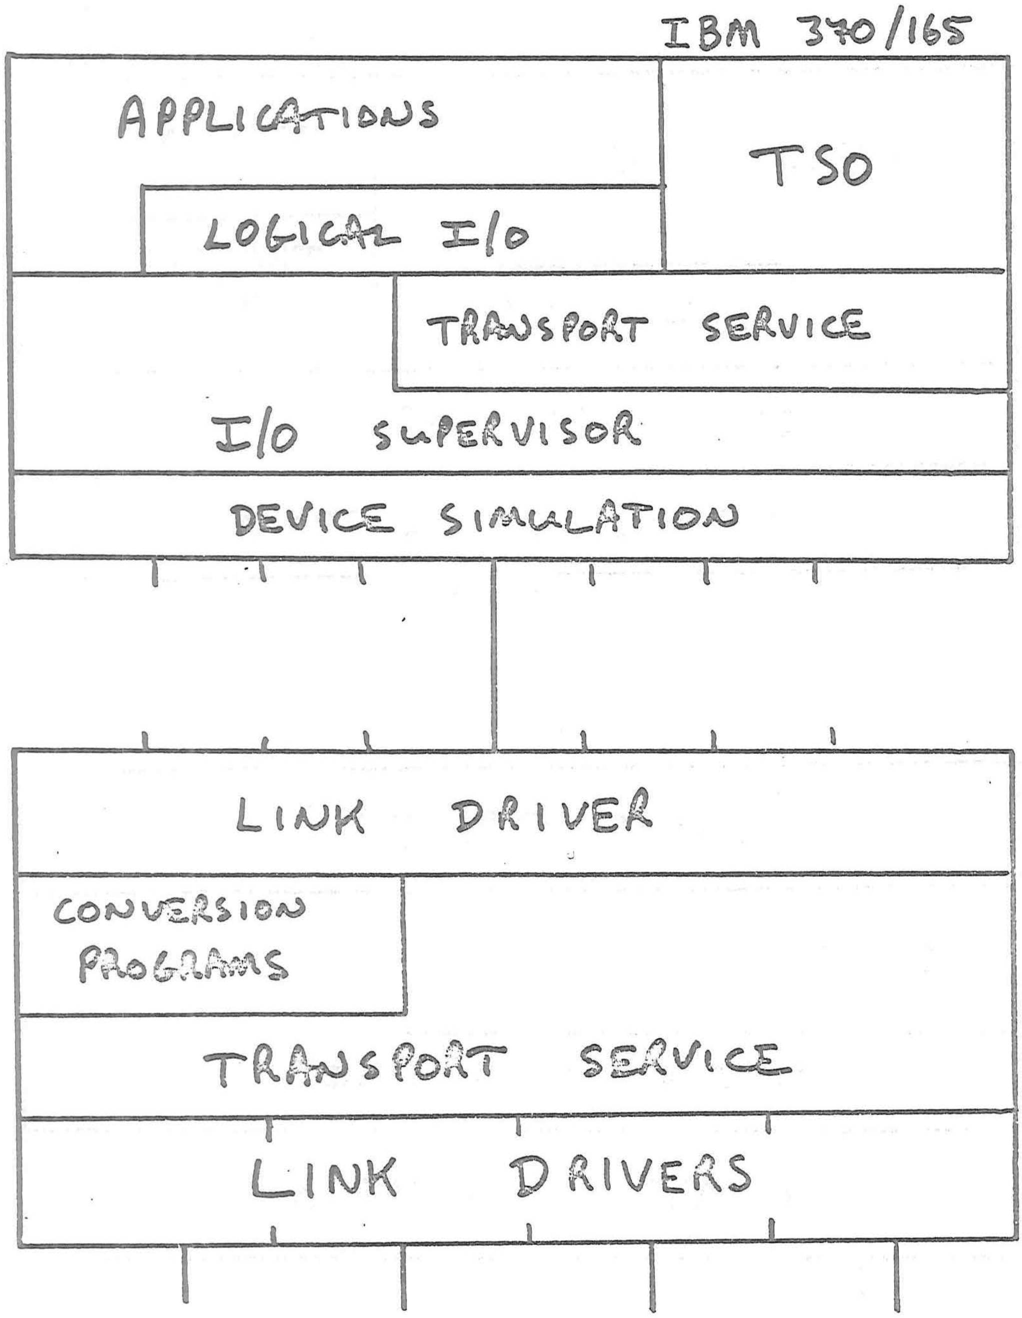

FIGURE 7. D.L. FRONT-END SYSTEM. PROPOSED

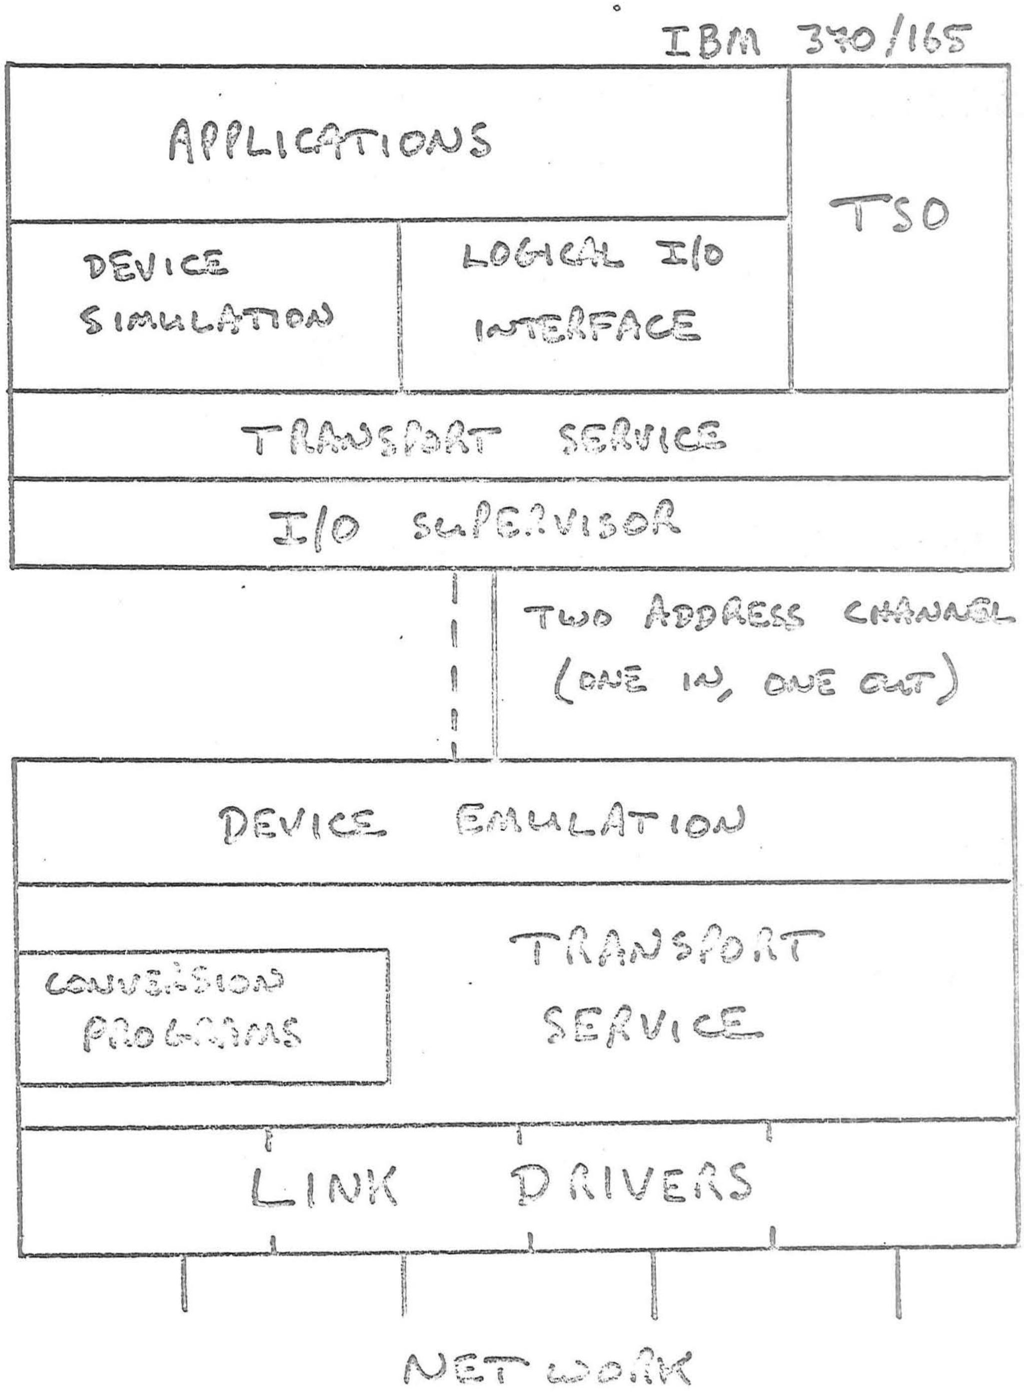

FIGURE 8. DATA RATES THROUGH FRONT-END

AVERAGE
- OVER 2 HOURS (IN 1977)
- FRIDAY 14.00 – 16.00 HOURS
- 35 ACTIVE TSO USERS
- 7 ACTIVE RJE
- SYSTEM CLASSED AS VERY BUSY

DATA RATE
- 8 BLOCKS/SEC IN EACH DIRECTION
- 200 BYTES/SEC IN EACH DIRECTION

PEAK
- CAUSED BY VERSATEC PLOTTER

DATA RATE (VERSATEC ONLY)
- 10 BLOCKS/SEC
- 5000 BYTES/SEC

Terminal Services in Networks

Dr. R. P. J. Winsborrow

Computer Science & Systems Division
AERE Harwell, Oxon, U.K.

Terminal Services In Networks

1. Introduction

The use of remote terminal access to computing facilities is common practice. Currently remote access to mainframes makes use of communications facilities provided by the PSTN and leased lines. It is natural to assume that the introduction of public packet switching communications networks on a national and trans-national basis will attract terminal communications traffic from the media now in use.

It is anticipated that packet switching networks (PSNs) may initially be used purely as replacements for traditional media for terminal communications. Terminal dialogues will, in these cases, continue to use manufacturer-specific protocols. However the use of networks can be expected to evolve to make use of the potential for communicating with a variety of computing facilities. This evolution can be helped, and to some extent anticipated, by the development of standard terminal handling protocols which permit ready interworking between terminals and services.

The aim of a standard terminal protocol is to enable communication between terminals and services irrespective of the specific model of terminal being used. It would be unreasonable to expect one single protocol to accommodate the whole range of terminal applications which can span the use of simple teletypewriters through to graphics. Instead it is expected that terminal protocols will be designed for particular application areas e.g. data entry.

The most immediate need is for a protocol which will permit communications between services and simple interactive terminals of the teletypewriter variety. These terminals constitute the bulk of the presently installed terminal population. The importance of the data entry applications has ensured that protocols for this area have also been given considerable attention.

In this paper we will describe the protocol developments within the CCITT and the standards organisations and look at the future prospects for standardisation.

2. CCITT Developments

Within CCITT the so-called 'triple-X' standard has been defined (X3, X28, X29). These recommendations define the facilities of a terminal concentrator for simple teletypewriter devices, the interface between the concentrator and a terminal user, and the protocol to be used for dialogues between a service and a terminal attached to the concentrator (figure 1). The concentrator or PAD (packet assembler/-disassembler) can be accessed via terminals using PSTN facilities. Communications between the PAD and an application make use of a switched virtual circuit provided by a PSN having a subscriber interface conforming to CCITT recommendation X25.

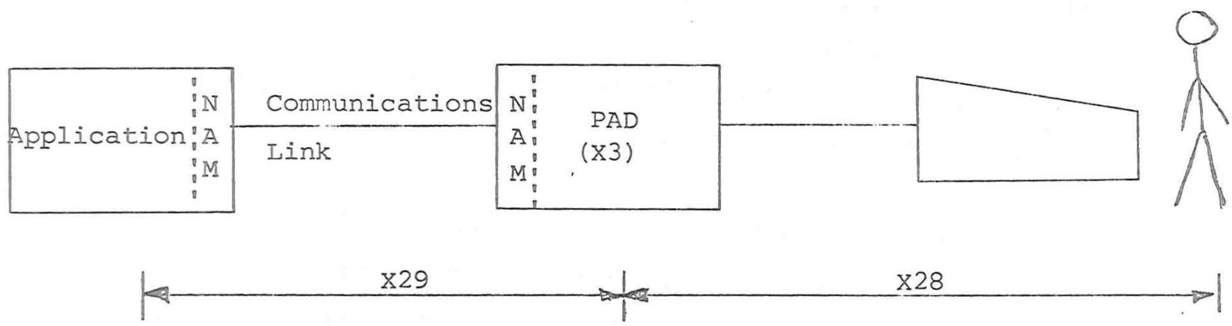

FIGURE 1 : TRIPLE X RECOMMENDATIONS

The triple-X recommendations represent the only internationally accepted set of protocols for interworking between simple asynchronous terminals and services on an X25 network. Therefore they are of some importance. The recommendations attempt to achieve device independence by specifying the use of International Alphabet No. 5. A measure of service adaption is provided by the capability to dynamically change data forwarding criteria in the PAD. The limitations of the recommendations include the need to use the Q-bit in X25 data packets and the inability of terminals to receive incoming calls.

A computing facility on an X25 network may well provide a PAD function for locally connected terminals wanting to access the network and an X29 acess method for remote network terminals to access available services. These access methods will be implemented using a basic X25 communications service. This set of recommendations probably represents the simplest and most economical way of coupling teletypewriter devices to services via a packet switching network.

3. ISO Developments

The development of standard terminal handling protocols for different terminal applications is being actively pursued within ISO by SC16 of TC97. This standardisation activity is based on work carried out by computer manufacturers and the research network community (notably ARPA, EIN and CYCLADES).

In contrast with the triple-X recommendations the ISO terminal protocols will be based on a communications service offered by a network-independent transport service rather than a specific type of network design. The facilities offered by a communications service of this type have been already described in this workshop.

Terminal handling protocols are being developed according to the philosophy of the virtual terminal concept. A standard or virtual terminal is defined which embodies many of the functions found in real terminals. It provides a means of referencing functions in a unique

manner even though they may be implemented in different ways on real terminals. Associated with a virtual terminal model a virtual terminal protocol is defined which controls the operation of the terminal. As the virtual terminal is an abstract entity real terminals and services using the virtual terminal protocol must adapt their access methods to those of the virtual terminal.

With the diversity of real terminals it is impractical to encompass all terminal applications using one virtual terminal model. Current developments have identified scroll mode, page mode and forms mode (data entry) application areas with their associated terminal models and protocols.

All virtual terminal models have a common basic structure which comprises:

- a presentation device for information display
- an input device for entry of data
- a store for the storage of data received from local and remote sources
- a control unit which manages access to the store and network.

There may be additional auxiliary input and output devices for certain applications. The complexity of the data structure which may be realised in the store and represented on the presentation device largely characterises a particular terminal model.

A virtual terminal protocol assumes that application entities converse using messages formed from elementary protocol commands. Commands are required for dialogue control, text display, addressing, display rendition, data entry access control, status monitoring, auxiliary device control, form definition and management. Strings of commands which may comprise a valid message are defined by the application.

Standardisation activities within ISO are slow and lengthy procedures because of the need to gain the agreement of a large number of interested parties with diverse backgrounds and interests. Although SC16 activities are being progressed with a good deal of vigour it is unlikely that a virtual terminal protocol standard will be achieved in less than three to five years. As an interim measure network users requiring the facilities of a terminal handling protocol may use the triple-X recommendataions or one of the protocols developed by EURONET, EIN, CYCLADES, ARPA or INWG.

4. Comparison of CCITT and ISO Developments

The CCITT recommendations are intended to cater only for simple teletypewriter devices; thus is equivalent to scroll mode operation in virtual terminal parlance. When comparing them in this application area the only major difference is in the dynamic adjustment of parameters effecting the terminal dialogue. Virtual terminal protocols only permit parameter adjustment for dialogue control (an area neglected by the triple-X recommendations); X29 permits adjustment of parameters which control the terminal/PAD interface and message forwarding signals. These facilities provided by X29 are assumed to be a purely local matter by virtual terminal protocols. The other main area of difference is that X29 is tied to the use of X25 - based network whilst the ISO developments will make use of a network - independent transport service.

5. Future Developments

The pressure to develop effective standard terminal handling protocols is likely to increase. This will arise not only from the traditional area of remote terminal networks but also from the development of computer services such as mailboxes and teleconferencing. Terminal handling protocols are actively being considered within BSI in WG4 of DPS20. This working group is also contributing to ISO developments.

Implementation of the File Transfer Protocol

Dr. P. F. Linington
University of Cambridge

Implementation of the File Transfer Protocol

P.F. Linington

Even after a file transfer protocol has been specified, there remain many issues to be decided in planning an implementation; this paper discusses the major problems likely to arise in implementing a file transfer facility under an existing operating system.

1. The protocol

The file transfer protocol to be implemented is here assumed to be that proposed by the High Level Protocol Group (December 1977) which provides facilities for the transfer of single sequential files and the preliminary exchange of information about their properties. The protocol design assumes that there is a clearly identifiable initiative responsible for the transfer, and that the judgement as to the required properties of and constraints on the transfer are located with this initiative. Thus one side of the dialogue will be the initiator (called P) and the other the responder (called Q). Independently of this distinction, one side will be the sender of the filed data and the other the receiver.

The protocol can be subdivided into a number of phases. The first is an initialization phase, which begins when the initiator states his aims in a start file transfer command (SFT). The responder replies, stating whether the transfer is possible, and if so whether any of the stated conditions need to be modified. Depending on whether the transfer is able to proceed or not, the initiator then signals a change of phase, either into the data transfer phase or the termination phase.

In the data transfer phase there is available a wide set of transfer control, error recovery, record structuring, coding and error reporting facilities. However, the use of these facilities is one of the aspects negotiated during the initialization, so that transfers can take place between implementations of differing levels of sophistication by use of a lowest common denominator. At the end of the data transfer phase, the protocol passes into the termination phase.

Details of the protocol will not be reiterated here, but copies of the specification can be obtained from the author of this paper.

2. Elements of the Implementation

The nature of a file transfer implementation is largely determined by the interfaces it must make with its environment, particularly to the operating system under which it is executing. At the initiator, attention focuses on the user interface, which may either be direct or via system control facilities, while at the responder the major problems are those of the filestore interface. On both sides an interface is needed to the communication facility, and it is assumed that this will be expressed in a general form as a Transport Service Interface. The objects or services to which these interfaces give access are complex. If they are not provided in a convenient form by the operating system, the effort needed to provide them may well be much greater than that required to produce the file transfer program as such.

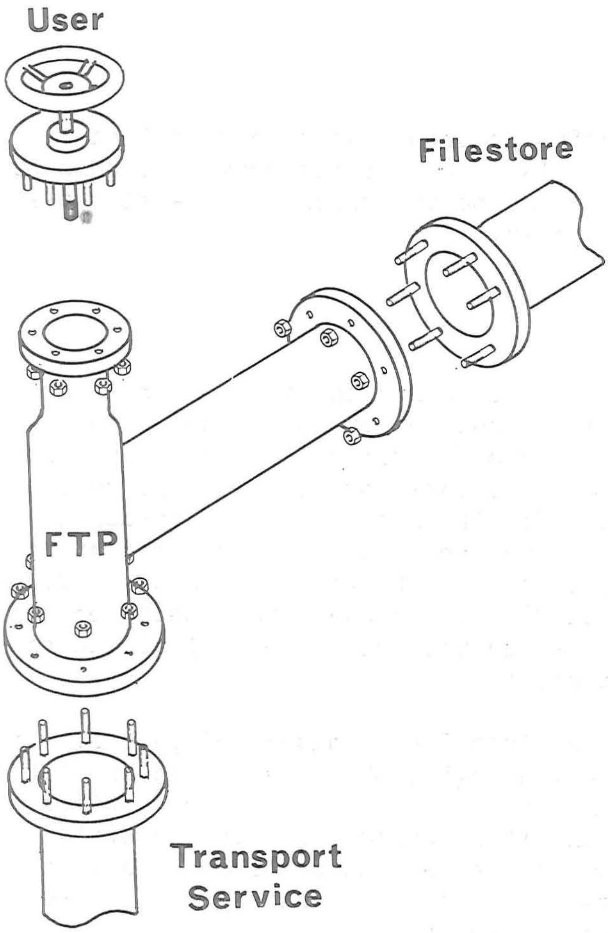

3. The User Interface

The user interface is almost entirely restricted to the initiator.

The complexity of the user interface can vary greatly between different user invironments. Perhaps the simplest situation arises when the user envokes the file transfer program directly as a utility; in this case there need be only one transfer per user command, and the syntax decoding and all local file accesses can follow existing local conventions. The user is available to receive error and status reports, and since the user has run the utility in his own protection environment, there are no additional problems of file security.

The situation is more complex when the file transfers are managed by a separate process, serving a queue of file transfer requests. Since the requests are handled by the one process, it is necessary to enter or simulate the effect of the user's protection regime, particularly when acting as receiver. The requirements for error and status reporting are also much more complex, because the route back to the user is longer.

Perhaps the most complex interfacing requirements arise if the file transfers are generated by the system as a result of requests expressed within a general job control or command and response language. In such cases the file transfer implementation must be integrated into the operating system to a substantial degree; depending on the nature of the operating system this may in itself be a very considerable project.

There will normally be very little direct user interfacing at the responder implementation, but data privacy legislation may specify user notification when a file is transferred.

4. The Filestore Interface

For the filestore interface, the major problems lie with the responder, which must obtain file description information from the protocol exchanges; however, some of the problems also affect an initiator which takes its requests repeatedly from a queue. Firstly, the program needs to control the filestore directly, during execution. This will be a trivial requirement for many systems, where the operating system provides a fully programmable interface, but in some systems file allocations and connections can only be made by the system when the program environment is created. In such systems much work may be involved in providing the necessary interfaces.

Secondly, the identity of the calling user is derived from the protocol, and so the program must be able to create the required protection or authorisation environment after it has started to execute. In some systems it may be possible to produce the required effect by creating a new task or submitting a job, thus using the existing mechanics for environment creation, but such measures are likely to be cumbersome and inefficient. A more satisfactory solution would be the provision of a protected procedure which, given a user identity and token of authority, such as a password, would change the environment of the calling program, perhaps temporarily, to that of the quoted user.

5. The Transport Service

The provision of a transport service is a prerequisite of the implementation of the protocol. Since few systems are currently supplied with communication support of this power or generality, it will need to be provided, and so this area does not pose the constraints of unsuitable existing software. Most of the current transport proposals assume an 8-bit byte interface, however, and this may pose considerable problems to machines whose basic operational unit is less than 8-bits. Indeed the Transport Service cannot, in general, provide even character code translation, as parts of the protocol data will be binary coded, and so must not be translated.

At the responding site, there is the problem of attaching the responding process to a received call. A primitive system may operate by maintaining a stock of listening processes, but this limits the degree of parallelism that can be supported and occupies significant resources in the idle state. The alternative is to generate a new process as a result of transport service stimulation of the system whenever a call is received. Means for assigning a suitable running environment for the new process are required, as already discussed above.

6. Local Actions

Use of a file transfer mechanism will involve a number of local activities which will vary considerably depending on the nature of the underlying system. For instance, there may be requirements for traffic monitoring or logging, or for collection of performance or tuning statistics. One activity which is very likely to be required is a form of accounting, and it is worth noting that there are three ways in which a file transfer may generate accounting information. Firstly, the use of the transport service resources will incur communications charges. Secondly, the access to the filestore will consume computational and input/output resources. Finally the transfer may result in the creation of a permanent filed record, which will incur storage charges.

7. Information on implementations

In view of the problems discussed above, a would-be implementor may save a substantial amount of effort, if he can obtain information on existing implementations. Such information may allow the avoidance of unsuspected problems, and may even result in the acquisition of an already available program. It should also reduce the danger of private network protocol implementations growing away from the main stream of the standard, like the Galapagos Turtle, through lack of contact with the outside world.

To provide a suitable forum for these exchanges, a File Transfer Protocol Implementors Group has been established within the UK, open to all who are in the process of implementing the protocol. This group exists to discuss implementational problems and hopes to maintain a register of the details of known FTP implementations in sufficient detail to provide assistance to would be implementors.

8. Conclusions

The effort required to implement a file transfer protocol depends crucially on the environment in which it must operate. If all the required operating system support functions and interfaces are available the implementor is fortunate and a simple file transfer facility can speedily be constructed. If, however, any of the major interfaces required are not available in the desired form, the effort needed to provide the facility may be dominated by that needed to make up the deficiencies. The effort will, however, ease the implementation of other protocols which may be defined in the future.

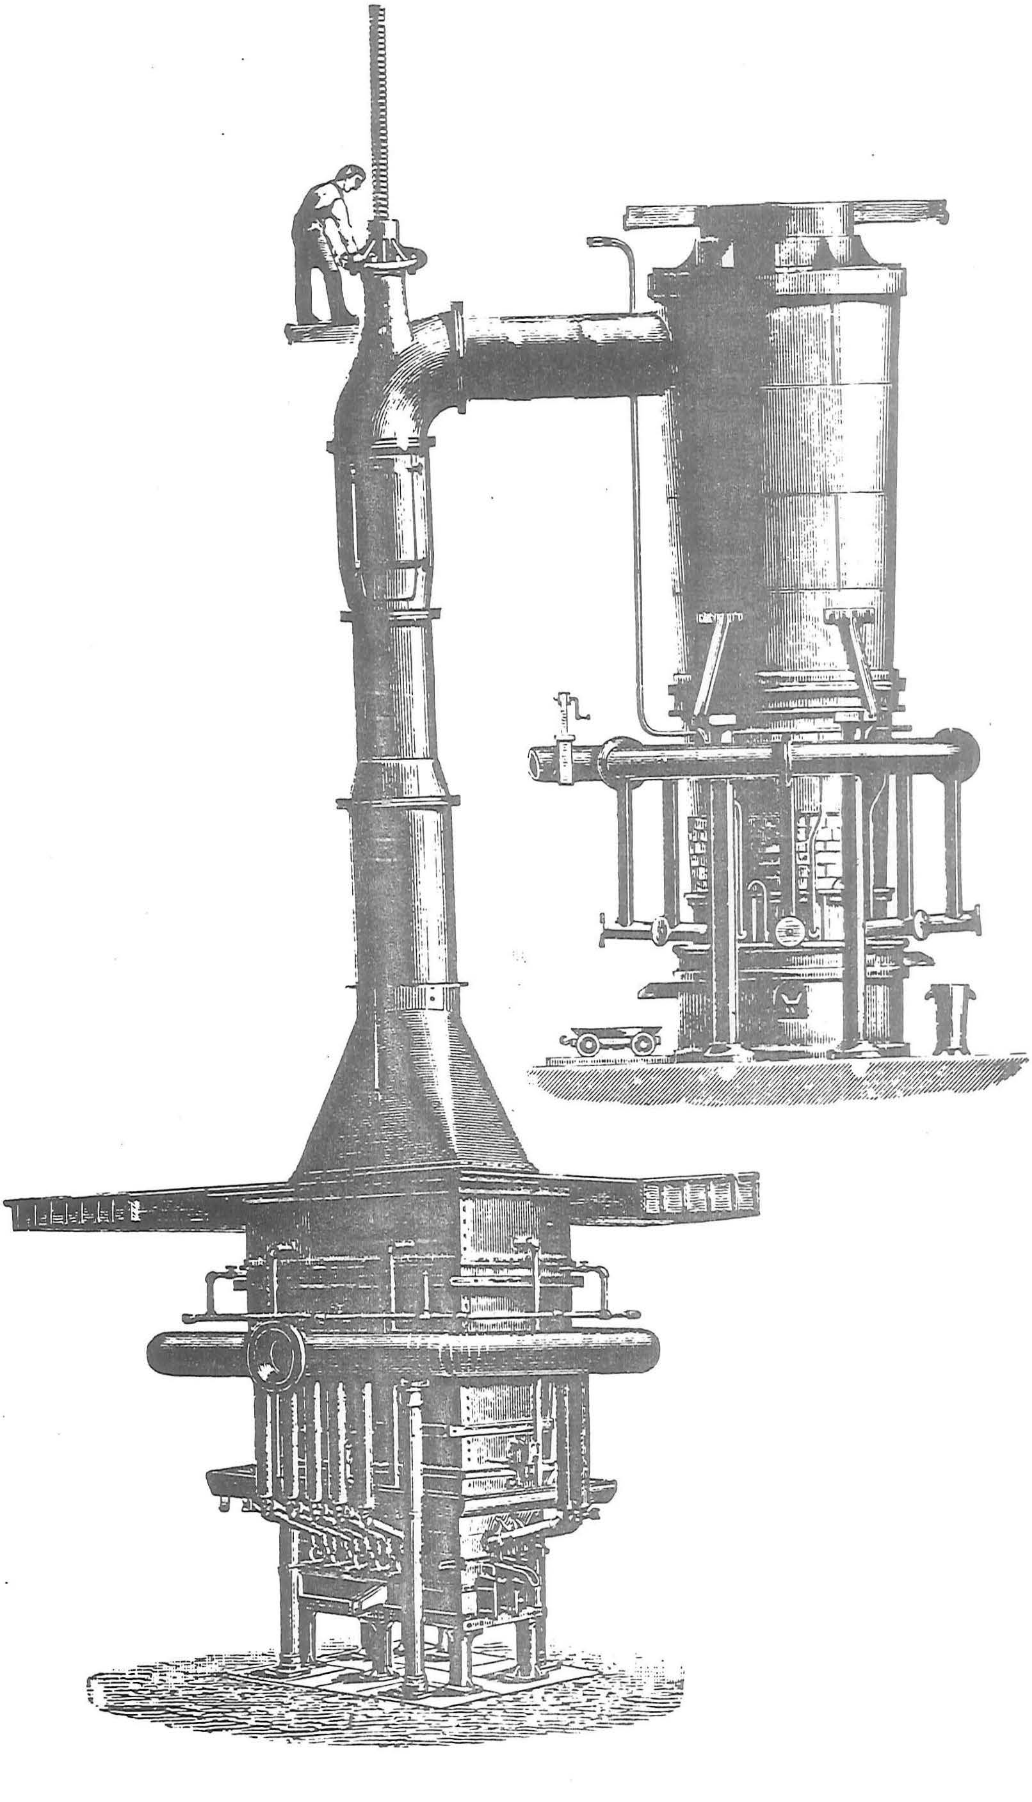

K.5

Requirements of a Job Transfer Protocol

Dr. M. A. Maconachie
Nottingham University

25th September 1978

REQUIREMENTS OF A JOB TRANSFER PROTOCOL

Dr. M. A. McConachie - Nottingham University

0. Summary

The topic to be discussed is Requirements of a Job Transfer Protocol, rather than just a protocol. This involves much wider issues concerning the environment in which the said JTP will be formulated and subsequently exist.

Thus the Requirements are reviewed in two ways -

a) those of the Macro Environment for the Design, Production, Service, Maintenance and Development of JTP.

b) those of the Micro Environment for the Host System, Network Services, User interfaces, together with supporting and component protocols.

Attention is mainly focussed on the possible provision of a single standardized JTP (since we appear to have several different ones already).

The following sections provide a little background to the 'micro environment' aspects of the subject. These paragraphs are in fact the introduction to a document I submitted as input to EPSS SG2 HLP Working Group in early 1976, but they outline a possible technique for specifying JTP's which is now more widely known and might thus form the basis for further discussion.

JOB TRANSFER PROTOCOL

1. GENERAL PRINCIPLES

The Job Transfer Protocol is intended to provide a method for the transfer of work, in the form of complete jobs or tasks, across a network to a remote machine, and for the subsequent retrieval of any output files resulting from the execution of that work. However, to be generally useful it must do more than that; it must provide a means of interrogating the remote Host about the jobs, and perhaps offer a means of controlling their progress, and of selecting the order and way in which their output returns.

One important consideration is that of the diversity of scale and sophistication between remote and local implementations. On the one hand, a simple terminal may wish to send jobs to a large central system providing a wide range of information and controls; while on the other hand a busy and sophisticated workstation may require to use a Host providing only the simplest of RJE services.

Another, related consideration, is that of a potential for future development. It is desirable that the form of the protocol allows for both the Terminal (the user) and the Host to adapt to changing demands and develop independently and progressively. This will often be in the direction of increasing sophistication. Such changes should be within the scope of the then existing protocol; but more importantly, because of the way in which the protocols will be implemented by a diverse body of users and services, a change must avoid at all costs the need for simultaneous changes by others.

One further consideration: the need for a Job Transfer protocol exists now, and so do the operating systems which will provide the RJE service. There is no way in which major (or in some cases even minor) changes to the operating systems of the Hosts will be possible. The protocol must therefore permit the maximum use to be made of existing facilities.

The general principles of the Job Transfer Protocol may thus be summarised as follows:

a) It sets out a way in which known characteristics of a job may be formally identified and transmitted in a standard format between Terminal and Host and vice-versa.

b) It provides for information, about the job or jobs already submitted for execution, to be obtained by the Terminal, again in a standard format.

c) Where the details of the description of jobs can be changed (i.e. the Host permits it) then these changes can be indicated and transmitted in a standard message format.

d) Where information is not accessible or simply not available this is allowed for in the protocol and may be indicated in a predefined way.

e) Details of a temporary or esoteric nature can be accommodated by mutual agreement and in a standard predetermined way at the time of their use. ('Mutual agreement' means that the Host implementation can accept (store) the information and/or the Terminal wishes to use it).

It should perhaps be emphasised that the JTP does <u>not</u> represent a Job Control Language. The actual jobs will be defined in the JCL appropriate for the Host's operating system.

On the other hand, the protocol could sometimes allow access to information normally specified within many JCL's, for example deadlines and output routing information. In some cases, this information may be changed and perhaps could be submitted via JTP.

2. BASIC CONCEPTS AND TERMINOLOGY

2.1 Hosts and Terminals

In the following descriptions the names Host and Terminal will be used to describe the two sites involved in job transfer. A Host is any computer system capable of executing the job, either directly or under its control. A Terminal is the place from which the job originated, and in most cases will be remote from the Host. Generally, this source or origin of the job will need to have a measure of control over the job, and, as often as not, will be the place to where any output will be returned. The Terminal may be any size of computer in its own right, from a simple RJE terminal based on a small mini computer, through more sophisticated workstations with access to several peripherals (and local networks, disc storage, etc.), up to a large mainframe computer system capable of executing jobs itself.

2.2 Jobs and Jobmills

The fundamental problem in attempting to define a job transfer protocol suitable for use by any Host is that the Hosts all define a job in a different way. They use a wide variety of terminology to describe the job's status and progress while being executed.

To overcome this the JTP sets out to define the characteristics of a job in a standardised way, and to establish a single formal description of job execution facilities that can be used (simulated) by all Hosts supporting the protocol.

Thus, the execution of the Host's background workload takes place in the Jobmill. The Jobmill has three places where the jobs are kept while in the Host; the Input List (well/queue), the Execution List, and the Output List.

The existence of a job in the jobmill is represented by an Entry in one or other of these lists. The job may only be in one of the lists at any time. As the job starts running it will move (the entry will be transferred) from the Input to the Execution list; upon completion it will be placed in the Output list. Note that this formal description does not refer to the physical files that hold the JCL and data or input, or to the lineprinter or other output files. It is solely concerned with a job as an administrative entity in a list.

The control and manipulation of input and output data files is of course a fundamental part of job transfer and it will be detailed later.

2.3 Attributes

Within the jobmill a job is represented by the values of a number of descriptive parameters known as Attributes. The exact number and details of these attributes will vary from Host to Host, but quite a few will be similar in format and use on most systems. For instance, nearly all Hosts require some form of alphanumeric jobname identifier, and usually record the origin of the job, and the target site for the output. These three, and several other attribute values (e.g. priority or deadline, account name, etc.), form the Entry for that job in the joblists.

Each attribute has a code number, a specification, a default value, and current value. The code number is in the range $0 - 255_{10}$ and serves to name and identify the attribute. The code number has two elements:

i) group number
ii) member number (within group)

The groups have been set up to divide the attributes by function and purpose, for example, group 1 for 'identification' attributes; group 2 for 'routing' attributes, etc. Each group can have a maximum of 16 member attributes.

The specification of the attribute indicates the data format to be used for any values that the attribute has, and also the current status of that attribute, for example, whether or not it is in use, or if the value may be changed or reset by the operator.

2.4 Jobmill Access and Control

Jobs and associated input and output arrive and depart from the jobmill through the Job Reception facility. In practice, this is a process (or processes) which can support job (i.e. file) transfers across the network. A certain number of parameters describe and govern a transfer - these are known as the transfer control attributes. (For instance, the data codes, the maximum record size, etc.).

The gathering of information about the jobs entered in the Lists, the control of the jobs in the Jobmill, together with the supply of details of the current status of the system, are the functions of the Operator Control facility. In order to seek information, and perhaps adjust the attributes of the jobs in the Jobmill, the Terminal can establish a link with the Operator Control facility via a process which supports Jobmill Interrogation and Control transfers. The establishment of this link may be compared with 'logging in' (logon, signon, etc.). The interchange of information which follows is during a Session. From a logical point of view the Session can be compared with a Job in the sense that it has associated attributes which describe and control it. For instance, Identification attributes such as a session name, the terminal name, account name, etc. have direct analogues in the Job Entry. There are several other groups of attributes, the values of which describe such things as the state of the system, the number of jobs in the Lists, the maximum number of simultaneous Jobfile Transfers allowed, etc. It is convenient to regard all of these as attributes of (i.e. available to) the Session.

The relationship of the various components of the Jobmill is illustrated in Fig.1. The Jobs are seen as sets of attribute values located in Entries in the three Lists. Note that this is a logical description of the JTP Jobmill and does not represent any particular Host system.

In passing, it is interesting to observe that a Filestore could be logically represented in an almost identical diagram. In Fig.2. the Files are seen as Entries in three Lists (Directories, Catalogues, etc.) The Lists again correspond to the possible 'places' that a file can be in. In this case the Interrogation and Control link would be used to inspect a directory, erase files, list them, archive them, etc, etc. It turns out that the control protocols that would be needed to perform these operations are almost identical to those described hereafter.

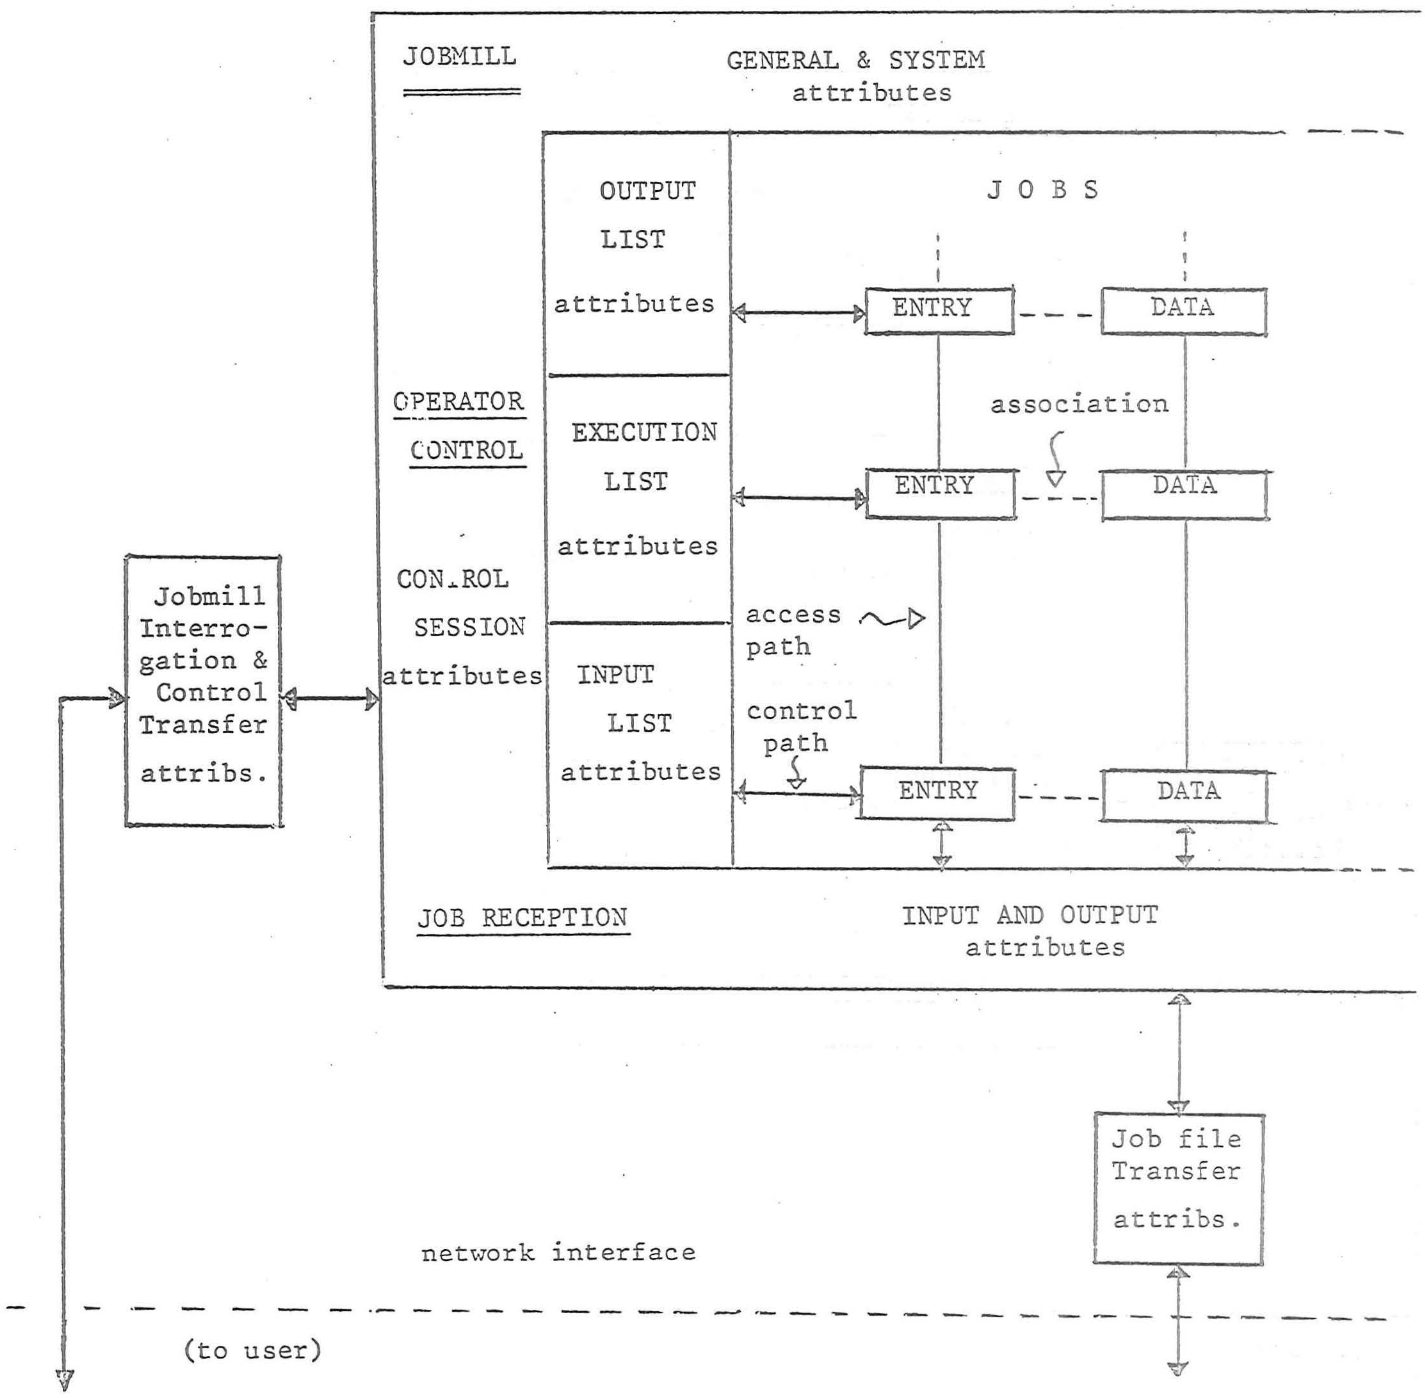

Fig. 1. A JOBMILL

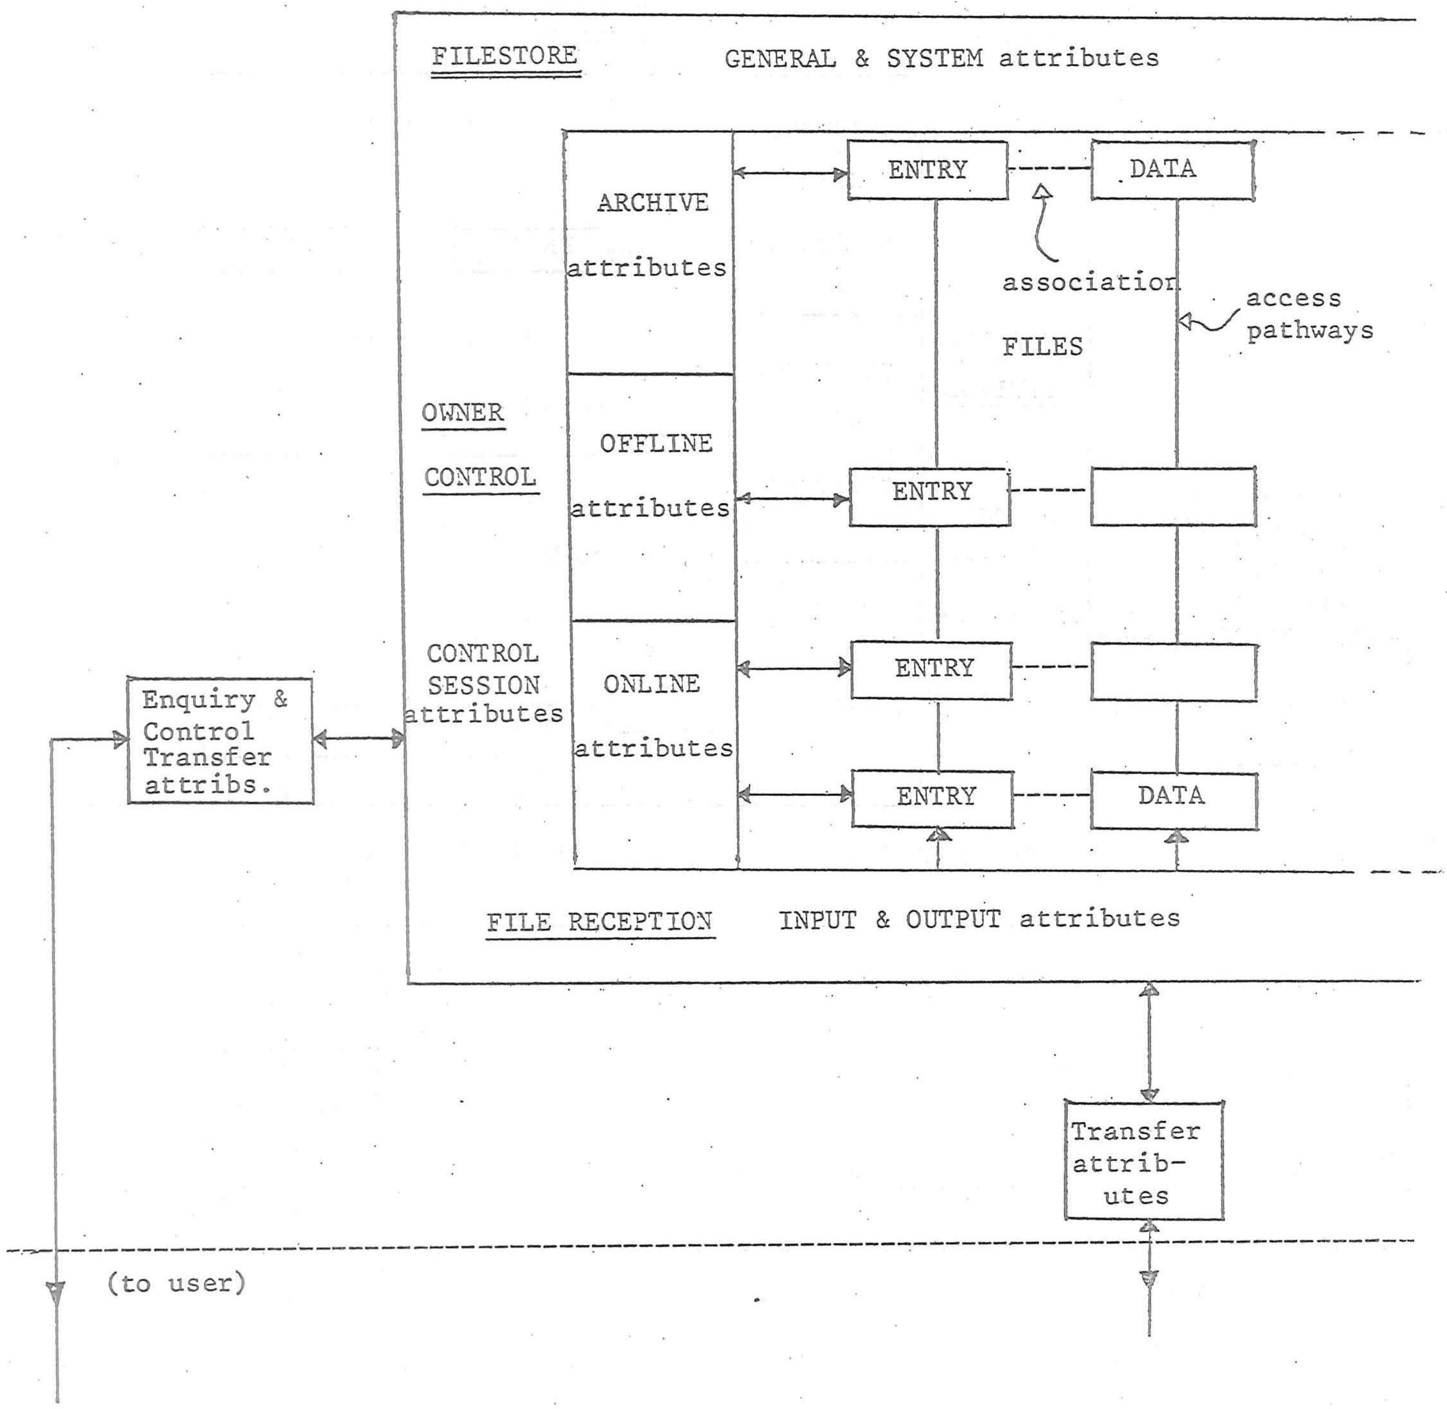

FIG. 2 - A FILESTORE

2.5 Component Protocols of JTP

From the preceding description of Jobmill it is seen that the JTP is composed of two sub-protocols :

a) Jobfile Transfer [JFTP]
b) Jobmill Interrogation and Control [JICP]

These two protocols are very similar and in both cases three levels of information interchange exist, namely :

Level 0 - Process initiation and control
Level 1 - Data flow control
Level 2 - The functional data of the process

a) Jobfile Transfer Protocol

At level 0 the remote Terminal is identified, and such matters as the data codes, compression, and primary direction of level 2 data transfer are established. The level 1 commands control the flow of data at level 2, and while they are in use the process will not recognise level 0 directives.

In this protocol level 2 data corresponds to the actual records of the job file being transferred.

b) Jobmill Interrogation and Control

The commands at level 0 establish the identity of the remote Terminal and set up the Session.

The nature of the process means that relatively small amounts of data are transferred and in most applications data flow control at level 1 will be largely unnecessary. However, in the interests of standardisation and to allow for future more sophisticated trans-network scheduling, they are included. Even in very primitive implementations the overhead is all but negligible.

The level 2 data in this protocol is made up of operator commands or directives, and the responses from the Host. Since the enquiries are about the values of the attributes of Jobs, & session etc. some of which are set up by level 0 commands, the format of a level 2 command is closely similar to those at the higher level.

The general nature of the command format will make it possible to specify other similar protocols (e.g. Filestore Interrogation) which use it.

2.6 Preliminary Notes on Implementation and Use

As mentioned previously the description of the Jobmill is of a formal, universally recognised entity into which all existing Host operating systems can be mapped. It is not intended that implementors create a replica in software or modify their current system to correspond to it. For instance, a system may not readily distinguish between jobs waiting to start and actually executing. In such a case any enquiry about the Input List would simply be given the reply that it was empty. In other cases the enquiry may refer to a piece of information (i.e. an attribute) that the Host does not keep or cannot provide access to. The protocols however provide for the Host to generate a standard value in the response,(in this case,a non-resettable 'null').

As a consequence of this approach the implementation can be very primitive initially and then built up as need dictates, and secondly the simple facilities of the Host can still be used by any Terminal, however sophisticated.

From a Terminal implementation standpoint, the choice of the level of sophistication rests locally. The JTP provides a way of gaining as much or as little information as required from the Host, and all in a standardised format. Thus the Terminal can have its own input/output scheduling system if required since the standardised information can be successfully used by programs.

Finally, although a Host may provide for control by the operator over the coming and going of jobfiles, this is not implicit in the protocols. Such notions as the setting up of input and output streams by the operator etc., are handled by the Terminal and not by the Host.

2.7 Summary

The JTP establishes the concept of a standard JOBMILL into which the implementor maps the existing job execution facilities of the local operating system. The Jobmill has JOB RECEPTION and OPERATOR CONTROL facilities.

Within the Jobmill JOBS are represented as JOB ENTRIES in three LISTS, INPUT, EXECUTION, and OUTPUT. Each Entry is a set of values for the

standardised ATTRIBUTES of a Job.

Operator enquiries and job administration are realised by the setting up of a SESSION. It is convenient to regard all the general parameters of the component parts of the Jobmill as being the ATTRIBUTES of a Session.

Jobs are input and output from the Jobmill by means of a process which supports the Jobfile Transfer sub-protocol (JFTP).

Operator enquiries and administrative control of the jobs are through commands and responses available by means of a process which supports the Jobmill Interrogation and Control sub-protocol (JICP).

Both protocols employ a standard three-level command structure; for Process control, Data Flow control, and for Functional Data of the process.

For JTP two sets of Attributes have been defined, namely

the JOB attributes → No.1

and the SESSION attributes → No.3

(The SESSION set includes the descriptive and control parameters of the Jobmill and its components, i.e. the Host's system).

The two lists below are extracts from more complete definitions of the two sets contained in Appendix A. They illustrate perhaps the more important member attributes of the sets.

(The meanings of 'format' and 'default values' are defined more fully in the following sections).

Note that both sets include attributes in the 'Transfer control' group which are provided for use only during the transmission of data (or commands and responses) into and out of the Jobmill.

JOB Attribute Set - Set No.1

Code No. (hex)	Name	Val.Resettable/ Fixed	Format	Default val.	Other Poss. vals. (e.g.)
[Transfer] Control					
00	Attrib.set no.	F	Number	1	<none>
01	Direction of Data Flow	R	Number	<Take>	<Give>
02	Codes	R	Number	<IA5>	(bit pattern as in FTP)
03	Format effectors	R	Number	<do nothing>	(")
04	Max.Rec.size	R	Number	<252>	≤252 octets
05	Marker Ack. Interval	R	Number		<32K>
[Identification]					
10	Jobname	F	Alpha	<blank>	<as given>
11	Jobnumber	F	Alpha	< 0 >	<given>
12	Username	F	Alpha	<visitor>	<given>
13	Account	F	Alpha	<none>	<given>
[Routing]					
20	Source Terminal	F	Alpha	<blank>	<source>
21	Target Terminal	R	Alpha	<source>	<as given>
[Control and Status]					
31	List (i.e. place)	F	Number	<Input>	<Execution> <Output>
32	Status	R	Number	<Run>	<Held> <Abort> <Delete>
[Logging]					
50	Date of Input	F	(Date)	<of event>	<none>
51	Time of Input	F	(Time)	<of event>	<none>
[Scheduling]					
60	Priority for execution	R	Number	<low>	<given>
61	Priority for output	R	Number	<low>	<given>
62	Mode of output	R	Number	<printed>	<punched> etc.

SESSION Attribute Set - Set No.3

Code No.	Name	Val.Resettable/ Fixed	Format	Default val.	Other Poss. vals.(e.g.)
[Transfer control]					
00	Attrib.set no.	F	Number	3	<none>
01	Direction of Data Flow	R	Number	<Give>	<Both>
02	Codes	R	Number	<IA5>	<as FTP>
"	(etc)				
08	Reply iden. no.	R	Number	0	<given>
09	Rep .sub-file max.no.of recs.	R	Number	32K	< 32K >
0A	Rep .sub-file actual no.of recs.	F	Number	< n >	
0B	Rep .rec.seq.no.	F	Number	< n >	
0C	Rep.. route tag	R	Number	0	<given>
[Identification]					
10	Session name	F	Alpha	<blank>	<given>
11	Session number	F	Alpha	< 0 >	<given>
12	Terminal name	F	Alpha	<blank>	<given>
13	Account	F	Alpha	<none>	<given>
14	Password	R	Alpha	<none>	<given>
[Routing]					
20	Source default for jobs	R	Alpha	<source>	
21	Target default for jobs	R	Alpha	<source>	
[Control and Status]					
30	No. of jobs in all lists	F	Number	< n >	
31	System status	F	Number	< n >	
32	Monitor changes to o/p List	R	Number	<off>	<on>
33	Monitor messages to "op"	R	Number	<on>	<off>

SOME

REQUIREMENTS

OF A

JOB TRANSFER

PROTOCOL

I.E. A SINGLE STANDARD
FOR JOB TRANSFER

NETWORKSHOP 3 SEPT. 1978

* MACRO ENVIRONMENT

- Design
- Implementation
- Service
- Maintenance
- Development

* MICRO ENVIRONMENT

- GENERAL NOTES
- HOST (TARGET)
- NETWORK
- USER (SOURCE)
- COMPONENT & RELATED PROTOCOLS
- USER/OPERATING INTERFACES

MACRO ENVIRONMENT

* DESIGN

  . THE TASK

   . ESTABLISHMENT OF WG; CONVENOR; SEC.

    . MANDATE; MONEY: TARGET DATE.

     . NUMBERS; REGULARS; ENTHUSIASM; EXPERIENCE

      . JOD; WRITTEN CRITIQUE & INPUTS.

* IMPLEMENTATION

  . ALL SITES
    - ENDORSEMENT AND ACCEPTANCE OF DESIGN.
    - COORD. OF TIMESCALES, VERSIONS, SOURCES/TARGETS.

  . PER SITE
    - PROGRAMMING EFFORT.
    - AUTHORITY TO USE RESOURCES.
    - DEVELOPMENT TIME; HOST & NET.

* SERVICE

- DAY TO DAY ADMIN; MANAGEMENT INFO.
  STATISTICS, HISTORY,
  ACCOUNTS & REPORTS.

- THE ANSWER TO "WHERE IS MY JOB?"

- PAS & DOCUMENTATION FOR REMOTE SERVICE.

- OPERATOR & SUPERVISORY TRAINING & FACILS.

* MAINTENANCE

- ERROR & EXCEPTION REPORTING.

- PROCEDURES & FACILS FOR USERS
  & SUPPORT STAFF TO DISCUSS DIFFICULTIES.

- FEEDBACK TO DEVELOPMENT MECHANISM.

- BACK-UP SERVICE & TRAINED SUPPORT STAFF.

* DEVELOPMENT

- ONGOING FORUM FOR CO-ORDINATION & SPECIFICATION OF STANDARD.
- SUPPORT IN MONEY, PEOPLE & TIME.
- FEEDBACK & INTERACTION WITH INTERNATIONAL DEVELOPMENT.

MICRO ENVIRONMENT

* GENERAL REQUIREMENTS OF THE ACTUAL JTP

. MUST <u>INTERFACE</u> WITH EXISTING OP.SYSTEMS

1. EASILY - I.E. NOT V. DIFFERENT.
2. CHEAPLY - MIN CHARGES TO SYSTEM.

. MUST <u>SUIT</u> WIDELY DIFFERENT SCALES

1. MINI TERMINALS.
2. MAXI HOSTS.

. MUST <u>ALLOW</u> PROGRESSIVE INDEPENDENT DEVELOPMENT

1. SIMPLE INITIAL IMPLEMENTATION.
2. TOTAL RANGE COMPATIBILITY.
3. ASYNCHRONOUS ENHANCEMENT.

* HOST (TARGET)

REQUIREMENT → A MODEL

E.G. A STANDARD
CONCEPTUAL
JOBMILL

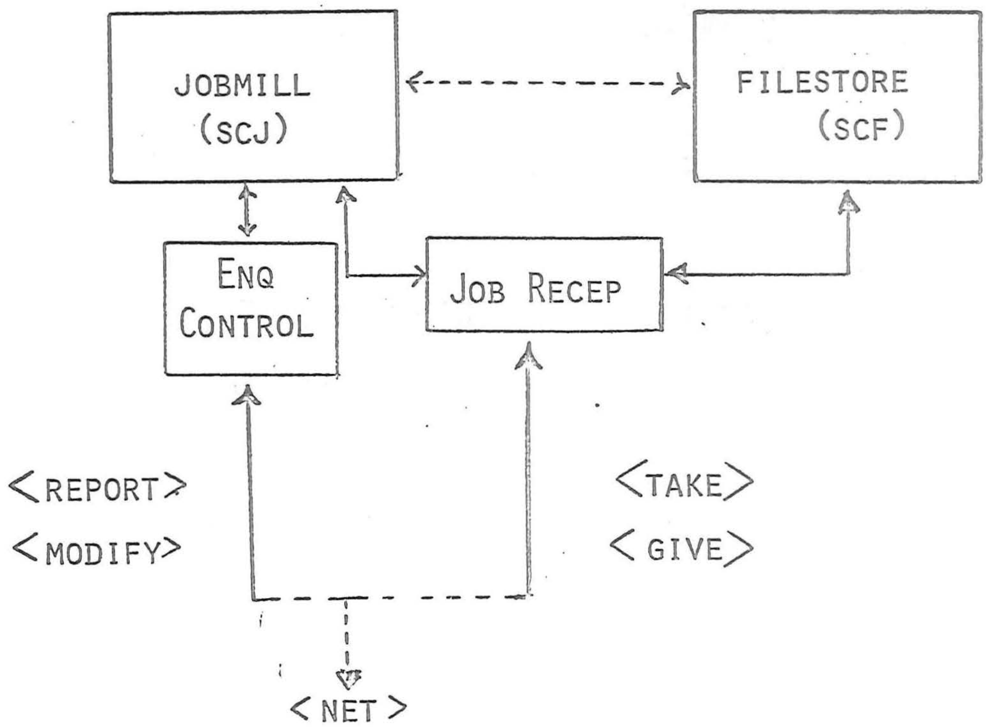

THUS AT THE HOST

JTP HAS THREE

COMPONENTS :

1. A STANDARD IMAGE → JOBMILL

2. A JOB FILE TRANSFER PROTOCOL → JFTP

3. A JOBMILL INTERROGATION & CONTROL PROTOCOL → JICP

* NETWORK

JOB TRANSFER PROTOCOLS SHOULD PLACE NO PARTICULAR REQUIREMENTS ON THE NET, EXCEPT FOR:

1. A BYTE STREAM IN EACH DIRECTION DIRECTION OVER THE LIAISON

2. SYNCHRONISATION (OR RESET)

* USER (SOURCE)

. NO MANDATORY FACILITIES IMPLIED BY JTP EXCEPT TO USE THE STANDARD ACCESS PROTOCOLS :

1) JFTP (E.G. FTP MINIMUM).

2) JICP (WHEN INFORMATION NEEDED).

* COMPONENT & RELATED PROTOCOLS

. JOB FILE TRANSFER

E.G. FTP WITH DIFFERENT OR
EXTENDED ATTRIBUTE SET.

. JOBMILL INTERROGATION & CONTROL

E.G. SAME STRUCTURE EXACTLY AS
FTP BUT FILE RECORDS ALSO
STRUCTURED LIKE COMMANDS
TO ALLOW MACHINE INTERPRETATION.

I.E. SEMI COMPILED NET JCL

* USER/OPERATING INTERFACES

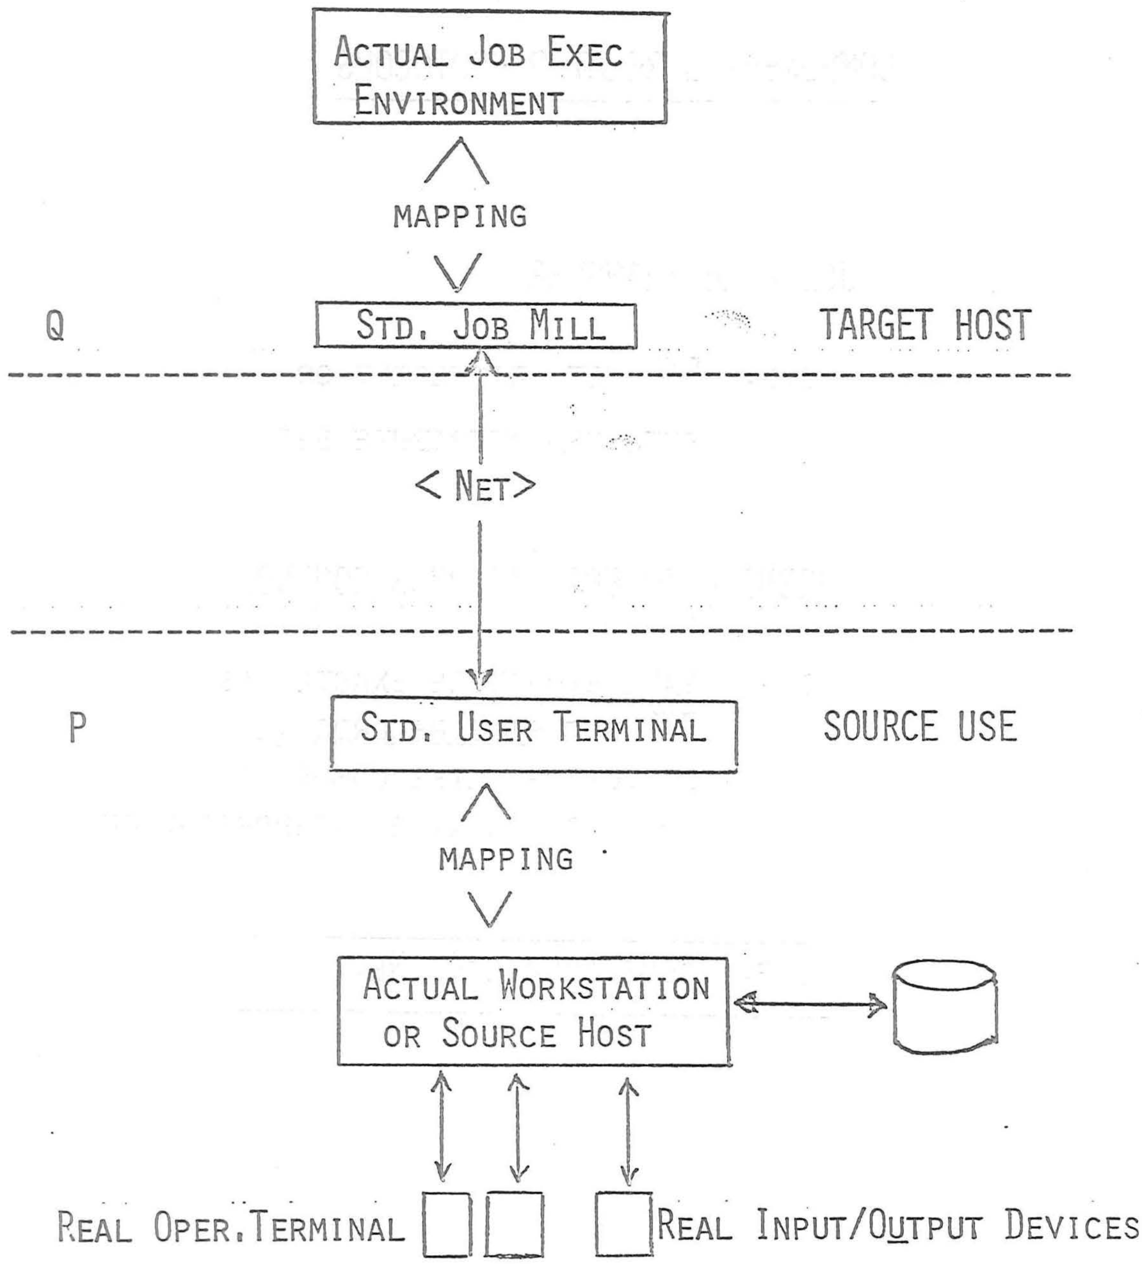

Report of the Work of the SWUCN Network Job Management Working Party

J. S. Thomas

South West Universities Regional Computer Centre
University of Bath

SOUTH WEST UNIVERSITES COMPUTER NETWORK

REGIONAL COMPUTER CENTRE

I U C C   N E T W O R K S H O P   3 ,   C A M B R I D G E

A report of the work of the SWUCN Network Job Management Working Party

Issued by
John Thomas

14 September 1978

0.0 CONTENTS

1.0 INTRODUCTION

The Network Job Management Working Party was established in the South West in early 1977; it has met ten times.

The main objective of the working party is to identify the basic Job Management facilities required in a network environment. The intention is that such facilities shall be incorporated into existing Job Control Languages e.g. by the use of macros. The definition of a complete 'Universal Job Control Language' including networking features is beyond the scope of the Working Party.

In its work to date the Working Party has adopted a 'top down' approach; however, conscious of the need to produce practical results we shall during the next six months attempt to map our ideas on to VME/B and Multics.

The author gratefully acknowledges the contribution of his colleagues: Alan Williams (Bristol), Kit Powell (Bath), Ian Parsons (Bristol), Rob Bradshaw (UCC), Wayne Clements (UWIST) and the previous chairman, Robin John.

2.0 NETWORK JOB MANAGEMENT FACILITIES

The following list identifies network job management user facilities.

2.1 Transfer of control within a job from one site to another.

This includes returnable and non returnable transfers in batch jobs as well as interactive connection to another site.

2.2 File handling

(i) Complete file transfer including routing to/from slow devices.
(ii) Remote file access at the block/record level.
(iii) Remote catalogue manipulation e.g. DELETE/CHECK.

2.3 Status facilities

A user should be able to determine the status of an individual site as well as the status of his job.

2.4 Job communication

(i) Remote job intervention e.g. suspension or termination.
(ii) Remote job communication via variables or messages.
(iii) Job to job communication

2.5 Resource allocation

A user must be able to distribute network wide the resources made available to him.

2.6 Implicit job transfer

Automatic high level scheduling.

3.0 TRANSFER OF CONTROL

3.1 A Network Job Environment (NJE) is proposed, that is a set of information associated with a job which is moved and updated as the job goes from one site to another. The NJE contains system-defined and user-defined variables. System-defined variables include jobname, username, account code, default print destination, status information and system-monitored result codes. User-defined variables can be defined, assigned and tested at any site.

Inside the NJE the variables are stored as manifestations of a simple machine-independent language i.e. in textual form.

This language would be block-structured and would initially handle string variables only. Manipulation of the variables at any site would be achieved by a particular realisation of this language suitably modified to suit the local job control environment.

3.2 Returnable transfer

Ideally the job control for a returnable transfer would look something like:

```
        site one job control
     AT 'site two' DO
        job control for site two
     OD
        site one job control (executed on return from site two)
```

The job control for site two and a copy of the NJE is transferred to site two. The local job control, journal and copy of the NJE (updated to indicate that the job has gone to site two) remain at site one.

The job continues execution at site two; when the OD is encountered the jobs output, NJE and journal return to site one, the NJE replacing the original copy and the journal being appended to that so far accumulated at site one.

In practice such embedded alien job control might cause severe syntactic problems for the local job control environment, particularly if, as envisaged, the alien job control can contain its own AT clauses. The basic facility would therefore be:-

```
     AT 'site two' DO 'file'
```

where 'file' contains the alien job control. The embedded version can then be provided as a luxury when feasible.

Nesting to any depth is inherent in the definition and constructs such as

```
     AT 'site two' DO
        AT 'site one' DO
        ...
        OD
     OD
```

must be catered for.

3.3 Attaching files

Individual files may be temporarily attached to a job and thus carried from one site to the next during a transfer of control. The filename and mode of access are specified as follows:-

AT 'site two' READING file one, file two APPENDING file three DO

The access mode defines how the files are to be treated:

READING	Move an existing file to site two; delete at site two when control returns.
APPENDING	File is to be created at site two, be returned and appended to an existing file at site one.
WRITING	Create the file at site two and return it to site one when control returns.
UPDATING	Move an existing file to site two, update it there and on return replace the original copy at site one.

The filenames 'fileone', 'filetwo' etc. must be machine independent and this is solved by defining them as NJE variables suitably linked to actual filenames at site one. The actual filenames used at site two need not be known by the user.

3.4 Non returnable transfer

The Job control for a non returnable transfer would look something like:-

```
   job control at site one
NGOTO 'site two' DO
   job control for site two
OD
```

The job control for the next site, the NJE and job journal so far produced are transferred to site two with any attached files. All knowledge of the job at site one is lost. As with any returnable transfers the user must explicitly identify any file to be moved with the job although in this case only READING is allowed.

3.5 The terminal interface

Connection to a remote site is achieved by a variant of the non-returnable transfer of control:

```
NGOTO 'site two'
```

The user is now connected to 'site two' and can LOGIN etc. Where the user is already logged in to a site and the machine presents a process-level interface (e.g. Multics, VME/B), the full transfer of control facilities defined above are available to him. When the destination machine also supports a process-level terminal interface, the alien job control can then be processed interactively rather than transferred en bloc, i.e.

```
LOGIN 'site one'                    interactive at site one
...                                      "        "   "    "
AT 'site two' READING 'file' DO
   ...                              interactive at site two
OD                                       "        "   "    "
...                                 interactive at site one
```

4.0 FILE HANDLING

4.1 File naming

A file naming scheme of the following form is proposed:-
<net file> ::= <remote catalogue name>. <local filename>
Standard catalogue names must be agreed but the local filename in this context refers to the name of the file as it would be catalogued in the catalogue at the remote site.

4.2 File transfer

Complete files can be moved either between file stores or between slow devices and filestores. Transfers may be queued on priority or FIFO queues either by users or processes. The completion of the transfer will be signalled by an event.

At a higher procedural level a process can be delayed until the event occurs, or the event may be re-ordered asynchronously by some suitable method, e.g. mailbox.

Sample job control command:-

COPY (net file 1, net file 2)

In the case of transfers involving slow devices e.g. printers and plotters a unique modified hierarchic addressing scheme is being developed to provide a means of naming objects in the network. (Reference NJMWP Discussion Paper 2.0 by Alan Williams) This scheme also deals with other entities casually referred to above e.g. sitenames, network file identifiers, etc.

4.3 File access

Processes should be able to connect a local I/O stream to a remote file and access individual records. The facility can be offered to terminal users if the process is run interactively.

4.3.1 Serial files (Formatted text files)

Data transfer between sites involves conversion to an agreed standard which needs to define character set, compression, format effectors, record length etc.

Read, write and append modes of access would be allowed and commands would be used to read or write the next record.

4.3.2 Direct access files (Unformatted binary files)

A self defining data scheme is required for the transfer of data from one machine to another. Access modes permitted would be read/write/append/update.

Any file access command understood by the remote operating system should be permitted (e.g. positioning) as well as read or write the next n bits.

4.4 Remote catalogue manipulation

Batch and terminal users should be able to effectively manipulate remote catalogues, e.g.

DELETE (net file)
CHECK (net file)
PROTECT (net file)

5.0 STATUS FACILITIES

5.1 Site status

A user or a process should be able to discover whether a site is in or out of the network and what facilities it is currently making available to users.

When used in network job control this allows conditional transfer of control constructs e.g.

IF 'site 1' THEN NGOTO 'site 1' DO etc.

At a more profound level, each intelligent entity in the proposed network would require a 'facilities list' accessible to the network job management software. For instance it is fully expected that some machines could not tolerate being the destination for a returnable transfer of control. The source machine must know this before transfer is attempted.

5.2 Job status

Job status information such as:-

JOB 'jobname' RUNNING AT 'site name'
-remote status details-

should be made available to the user. This information could be obtained by network job management software following through a chain of NJE's until it reaches the site where the job is currently executing.

6.0 JOB COMMUNICATION

6.1 Job intervention

The ability to suspend, restart or terminate a remote job is fundamental.

6.2 Job communication

A user should be able to communicate in general terms to, for example, change scheduling parameters such as job class or time limit that may be stored in the NJE. Changing job space variables can also be contemplated although exchange of large quantities of data between job and user is excluded; this is best handled by running the job interactively.

6.3 Job to job communication

This strays into the area of task activation and control, a topic being addressed by BSI/DPS20/WG2.

7.0 RESOURCE ALLOCATION AND IMPLICIT JOB TRANSFER

7.1 Resource allocation

The working party have been unable to devote much time to this topic; suffice it to say that the aim would be to allow users to spend their allocation of network resource units at any host to which they have access permission.

7.2 Implied job transfer

Certain types of job exist which do not need to run at a particular site i.e. the jobs are site or machine independent and the user is not concerned about where the job is actually run. If the user can indicate in some way to the network this property of his job then the network has the opportunity of deciding where to run it, with potential advantages in balancing work loads between machines. Application packages driven by macros are particularly susceptible to this treatment.

8.0 CONCLUSION

The paper describes the progress made to date by the SWUCN Network Job Management Working Party. Much detailed work still needs to be done but it is believed that the main elements of future Network Job Management have been identified and at least partly defined.

More detailed specifications for the facilities will be produced in due course; in particular implementations are planned for the Avon Multics system and SWURCC 2980 during the next eighteen months.

APPENDIX A

Investigation of other Relevant Work

Anxious not to re-invent the wheel, the Working Party has searched the literature for descriptions of other work on Network Job Management. A selected bibliography is given below. A number of other groups are doing theoretical work in the area, e.g. a group at ETH Zurich working on EIN, but we could find no evidence of a practical scheme having been implemented and in regular use by end-users. Some specific areas have been implemented, such as the ARPANET Datacomputer, which provides a network filestore. Further investigations are being made in the hope of discovering other areas.

One result of these investigations is that no-one yet seems to have solved all the problems of Network Job Management and there is no scheme in operation anywhere which encompasses a wide range of facilities. Against this background, the networking features of SWUCN Tasking JCL are seen as being quite advanced and worthy of being publicised more.

Selected Bibliography on Network Management

1. BCS-JCL Journal of Development (available from Ian Parsons)

2. Job control in heterogenous computer network, Schicker, Bacchi, Duenki; EIN Project.

3. Command language design for networks of processors, Newman, Fitzhugh; Loughborough University.

4. Command languages and heterogenous networks, Chupin; CII, Grenoble Scientific Centre.

5. Past and present solutions to the problems of network job control, Rayner; NPL.

APPENDIX B

Terms of reference of thé Working Party

The following terms of reference were established at the first meeting of the Working Party:-

To investigate the basic Job Management facilities required in a network environment, taking into consideration the following points:

(a) the development of on-campus networks, involving small, medium and large machines, and the connection from them into regional and national network;
(b) the advent of heterogeneity in mainframe systems;
(c) the needs of both batch and terminal users;
(d) the ability to implement consistent subsets of the full range of facilities;
(e) the impact on the network in terms of general loading and response times;
(f) the problems of implementing the facilities on current and future machines.

Conclusions and Future Activity

M. B. Williams
Network Unit of the Computer Board and Research Councils

CONCLUSIONS AND FUTURE ACTIVITY

Chairman: Dr R.A.Rosner

Speaker: M.B. Williams

Summary of Networkshops: what needs doing

It is to be regretted that the Post Office have not yet reached a definite conclusion about PSS nor have announced a customer engineering facility that would allow a detailed and fast dialogue over technical details. The addressing scheme is critical and the user community could be getting on with the provision of suggestions in this area. (Roland Rosner would like to see the main features of this in connection with PSS outlined by February 1979.)

Who is a potential user of PSS? The tariff issue is crucial in determining the answer to this question. The Network Unit may send round a questionnaire on this shortly. This could involve questions about network services in general, particularly in connection with FTP and JTP. Some useful results may be gleaned from this. The Network Unit will discuss testing arrangements with the Post Office.

Campus networking

A broader presentation of this subject might have been useful. Campus networking should be inherent in the computer system provided at any site and those people involved in systems and funding at such sites should appreciate this. This is a potential area of future activity for the Network Unit.

Mainframes

Network adaptation needs further consideration on most ranges of mainframe computers. A DEC 10 contract has been placed; discussion (but little progress) on 2900's continues and work on 1900's is about to begin. The Network Unit will consult interested ICL sites about the last two machine series. The Computer Board has circulated a letter to all main manufacturers giving operational requirements for networking to be effective on all machines delivered after 31 December 1980. It is possible that if a PDP/11 networking group were set up it might generate considerable interest. A list of people involved would be useful - Paul Kummer is the best person to contact for those wishing to be on such a list.

Transport Service

There seems to be uncertainty over what action is needed. A solution to this question is needed fairly quickly. Keith Bartlett (NPL) would like to arrange a half-day seminar on the Transport Service.

Terminal Handling

There is a need for a summary comment on what decisions should be taken - for example do we need XXX and another protocol? Or if not, then what?

FTP

The FTP working group is currently without a parent body, although it is partly under the wing of the Post Office.

JTP

A JTP working group is needed. It would be extremely useful if one person were to generate a first-iteration document for provoking discussion. In the longer term an FTP-type document could perhaps be produced by a group similar to the FTP group. A possible basis for this might be the High Level Protocol group with perhaps a slightly changed membership. This would produce funding problems though.

Networkshop 4

Networkshop 4 is planned to take place at York in the Easter vacation of 1979. It is perhaps time to start considering an ongoing organisation for Networkshops.

Appendix 1

Delegates

D. Johnson
Manchester University
Oxford Road
Manchester M13 9PL

C.J. Kennington
University College London
17-19 Gordon Street
London WC1H 0AH

K. Knightson
Post Office Telecommunications
River Plate House
Finsbury Circus
London EC2M 7LY

Dr. P.S. Kummer
Science Research Council
Daresbury Laboratory
Daresbury, Warrington WA4 4AD
Cheshire

Dr. J. Larmouth
University of Salford Computing
Laboratory
The Crescent
Salford M5 4WT

Dr. P.F. Linington
University of Cambridge Computer
Laboratory
Corn Exchange Street
Cambridge

Dr. J.A. Linn
University of Aberdeen
Computing Centre
150 Don Street
Aberdeen AB2 1 SQ
Scotland

G.W. Litchfield
Oxford University Computing Service
13 Banbury Road
Oxford

N.P.C. Lynch
University of Reading
Computer Centre
Whiteknights
Reading, Berkshire RG6 ZAX

Dr. A.J. Mascall
University of Newcastle Upon Tyne
(NUMAC)
Computing Laboratory
Claremont Tower
Claremont Road
Newcastle upon Tyne NE1 7RU

P. Mason
University of Sheffield Computing
Services
Sheffield S10 2TN

Dr. M.A. McMconachie
Nottingham University
Cripps Computing Centre
Nottingham NG7 2RD

Dr. J.A. McGinnety
National Environmental Research
Council
Alhambra House
27-33 Charing Cross Road
London WC2H 0AX

B.W. Morgan
South West Universities Regional
Computer Centre
University of Bath
Claverton Down
Bath

C. Morris
Network Unit of the Computer Board
and Research Councils
c/o Rutherford Laboratory
Chilton, Didcot
Oxon OX11 0QX

J.E. Morris
University of Leicester Computer
Laboratory
University Road
Leicester

P. Morrison
Post Office Telecommunications
Lutyens House
1-6 Finsbury Circus
London EC2M 7LY

D. O'Sullivan
Queen's University
Belfast
Northern Ireland BT7 1NN

B.J. O'Mahoney
University of Essex
Computing Centre
Wivenhoe Park
Colchester CO4 3SQ
Essex

D.M. Pick
Queen Mary College Computer Centre
Mile End Road
London E1 4NS

C.J. Powell
University of Bath
Bath BA2 7AY

R.G. Powell
Centre for Computer Studies
The University
Leeds LS2 9JT

J.A. Prentice
Loughborough University Computer Centre
Ashby Road
Loughborough, Leicestershire

J.E.H. Quilliam
Computing Unit, University of Surrey
Guildford
Surrey GU2 5XH

P. Ramm
University of Bradford
Computing Laboratory
Richmond Road
Bradford BD7 1DP

Dr. J.D. Rice
University of Liverpool Computer Laboratory
P.O. Box 147
Liverpool L69 3BX

I.M. Richmond
Rothamsted Experimental Station
Harpenden, Hertfordshire

C.C. Ritchie
Brunel University
Uxbridge, Middlesex UB8 3PH

Dr. R.A. Rosner
Network Unit of the Computer Board and Research Councils
c/o Rutherford Laboratory
Chilton, Didcot
Oxon OX11 0QX

A.J. Scolley
University of Stirling
Stirling
Scotland FK9 4LA

J.D. Service
The University of York
Department of Computer Science
Heslington
York YO1 5DD
North Yorkshire

D.J. Spridgeon
University of Hull
The Computer Centre
Cottingham Road
Hull HU6 7RX

R. Thirlby
Loughborough University Computer Centre
Ashby Road
Loughborough
Leicestershire

J.S. Thomas
S.W.U.R.C.C.
University of Bath
Claverton Down
Bath BA2 7AY

David Turner
University of Liverpool
P.O. Box 147
Liverpool L69 3BX

R. Vivian
University of Birmingham
The Computer Centre
P.O. Box 363
B15 2TT

Dr. D. Wanless
University of Keele
Keele
Staffordshire ST5 5BG

Anthony West
Computer Systems Laboratory
Queen Mary College
Mile End Road
London E1 4NS

Professor M.V. Wilkes
University of Cambridge Computer Laboratory
Corn Exchange Street
Cambridge CB2 3QG

A.H. Williams
Avon Universities Computer Centre
School of Mathematics
University Walk
Bristol BS8 1TW

M.B. Williams
Network Unit of the Computer Board and Research Councils
c/o Rutherford Laboratory
Chilton, Didcot
Oxon OX11 0QX

Dr. R.P.J. Winsborrow
AERE, Harwell
Building 7.12
AERE, Harwell
Oxon OX11 0RA

J.W. Anderson
Lancaster
Lancaster University Computer Services
Bailrigg
Lancaster LA1 4YW

V.G. Aswani
University of Manchester Regional Computer Centre
Oxford Road
Manchester M13 9PL

R.D. Atherley
New University of Ulster
Coleraine
Co. Londonderry
Northern Ireland

Peter Barry
University of Glasgow Computing Services
Glasgow G12 8QQ

K. Bartlett
National Physical Laboratory
Division of Numerical Analysis and Computer Science
Teddington
Middlesex TW11 0LW

Dr. R. G.Blake
Essex University Computing Centre
Wivenhoe Park
Colchester CO4 3SQ, Essex

R. Bland
The Hatfield Polytechnic
School of Information Sciences
P.O. Box 109, College Lane
Hatfield, Herts. AL10 9AB

R.G. Bradshaw
University College, Cardiff
Computing Centre
39 Park Place
Cardiff CF1 3BB

P.R. Brady
University College of Swansea
Singleton Park
Swansea SA2 8PP

J.P. Brandon
Queen Mary College
Computer Centre
Mile End Road
London E1 4NS

Dr. P.E. Bryant
Science Research Council
Atlas Computing Division
Rutherford Laboratory
Chilton, Didcot
Oxon OX11 0QX

John Burren
Rutherford Laboratory
Science Research Council
Chilton, Didcot
Oxon OX11 0QX

C. Bye
University of Dundee
Computing Centre
Park Place
Dundee DD1 4HN

J.G. Byrne
Department of Computer Science
Trinity College
Dublin 2, Eire

A.M. Chambers
University of Bristol Computer Centre
University Walk
Bristol BS8 1TW

A. S. Chandler
Computer Aided Design Centre
Madingley Road
Cambridge

R. Chisholm
Edinburgh Regional Computing Centre
The Kings Buildings
Mayfield Road
Edinburgh EH9 3JZ

T.B.G. Clark
University of Warwick Computer Unit
Coventry CV4 7AL

S.R.M. Clelland
Heriot-Watt University
Department of Physics
Riccarton
Currie
Midlothian EH14 4A5

N.H.R. Cooper
University of Sussex Computing Centre
Falmer
Brighton BN1 9QH
Sussex

Mr.I.N. Dallas
University of Kent
Computing Laboratory
Canterbury
Kent CT2 7NF

Dr. H. Davies
Computer Unit
University of Exeter
Exeter, Devon

P.M. Dewar
University of St. Andrews Computing Laboratory
North Haugh
St. Andrews
Fife KY16 9SX, Scotland

J. Down
University of London Computer Centre
20 Guilford Street
London WC1N 1DZ

G.R. Eadie
Computer Unit
Durham University
Durham

P.M. Girard
Rutherford Laboratory
Chilton
Didcot OX11 0QX

H. Gluck
Imperial College Computer Centre
Royal School of Mines
Prince Consort Road
London SW7

M.J.T. Guy
University of Cambridge Computer Laboratory
Corn Exchange Street
Cambridge

P. Hare
Open University
Walton Hall
Milton Keynes, MK7 6AA

P.S. Harrison
Nottingham University
Crpps Computing Centre
Nottingham NG7 2RD

Dr. D.F. Hartley
University of Cambridge Computer Laboratory
Corn Exchange Street
Cambridge CB2 3QG

Dr. W.D. Hay
Edinburgh Regional Computing Centre
The King's Buildings
Mayfield Road
Edinburgh EH9 3JZ
Scotland

Dr. K.S. Heard
AERE, Harwell
CSSD, H7-12
AERE, Harwell
Oxon OX11 0RA

P. L. Higginson
University College London
Statistics and Computer Science Department
Gower Street, London WC1E 6TB

A. Hinchley
University College London
Statistics and Computer Science Department
Gower Street
London WC1E 6TB

A.D. Holt
The City University (Computer Unit)
St. John Street
London EC1

B. Hook
Southampton University Computing Service
The University
Southampton SO9 5NH

A. Hopper
University of Cambridge Computer Laboratory
Corn Exchange Street
Cambridge CB2 3QG

J.B. Jamieson
Strathclyde University
204 George Street
Glasgow G1 1XW

David Jennings
UWIST
King Edward VII Avenue
Cardiff, S. Wales